Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

September 10, 2016

Dr. Rupak Biswas – Project Manager
NASA Ames Research Center, Moffett Field, CA
Rupak.Biswas@nasa.gov
(650) 604-4411

Facilities Team Works with Contractor to Meet Major MSF Milestone

- HECC Facilities engineers, in coordination with NASA Ames' Center Operations (Code J) staff, oversaw the fulfillment of a major milestone for the prototype Modular Supercomputing Facility (MSF) pilot project — the completion of the concrete foundation.
- The initial MSF module will house 16 SGI racks with Intel E5-2680v4 (Broadwell) processors.
- If the prototype proves successful, the next step will be to install a second module on the foundation.
- The foundation is 44-ft. by 67-ft., is 16-in. thick, and required 180 cubic yards (20 truckloads) of concrete.
- In addition to completing the foundation, conduits for the transformer, switchgear, and network were set in place prior to the pour.

Mission Impact: By working closely with the Center Operations staff and the construction contractor, the HECC Modular Supercomputing Facilities (MSF) project team ensured the successful preparation for and execution of the concrete pad pour for the MSF.

From left to right: Chris Tanner, Bill Thigpen, Karen Petraska (NASA Headquarters), Piyush Mehrotra, and Chris Buchanan standing on the newly completed Modular Supercomputing Facility foundation.

POCs: Bill Thigpen, william.w.thigpen@nasa.gov, (650) 604-1061, NASA Advanced Supercomputing (NAS) Division; Chris Tanner, christopher.tanner@nasa.gov, (650) 604-6754, NAS Division, CSRA LLC

Successful Suspend/Resume Process Used for HECC Systems Maintenance Activities

- HECC engineers recently used a suspend/ resume process on HECC systems to successfully perform important maintenance activities.
- This was the first time the Systems team used a suspend/resume method instead of a traditional system downtime. This method reduces the impact on users by avoiding draining jobs from the queues several days ahead of dedicated time.
- Maintenance activities included patching infrastructure systems, updating firmware, changing InfiniBand cabling, and other work.
- In the past, suspend/resume was only used for adding or removing compute nodes from Pleiades.
- The process is not risk free, as suspended jobs can fail to restart for a variety of reasons. HECC engineers continue to evaluate methods to minimize job losses and improve the robustness of the process.

Mission Impact: By utilizing the new suspend/resume functionality, HECC minimized interruption to users while performing required maintenance activities.

Regular maintenance on the Pleiades supercomputer and other HECC systems provides a stable, well-performing environment for NASA users.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing (NAS) Division, Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NAS Division, CSRA, LLC

Version 1 of myNAS Website Released to HECC Management

- The website, myNAS (version 1), developed as a collaboration between the HECC Application Performance & Productivity team and the Tools team, was released to HECC management. The website provides near-real-time insights into the ways missions, projects, and individual scientists are using HECC resources.
- The site provides high-level and detailed information on running and queued jobs, SBU rates, and allocation usage through a series of interactive tables and charts. The charts contain information including: SBU usage rates by mission, waiting jobs, mission target allocations vs. usage, and the number of Pleiades systems it would take to handle the current queued jobs.
- Numerous scripts deliver data to the site from qstat, Acct Query, and API calls every 10 minutes. Users can roll over charts to drill down for detailed information.
- Future releases of the myNAS website are planned for principal investigators (PIs) and individual users. Users will be able to monitor the status of their jobs and allocations, while PIs will be able to do the same for all members of their groups.

Mission Impact: The myNAS website provides management with a tool to quickly see the current status of HECC systems' jobs.

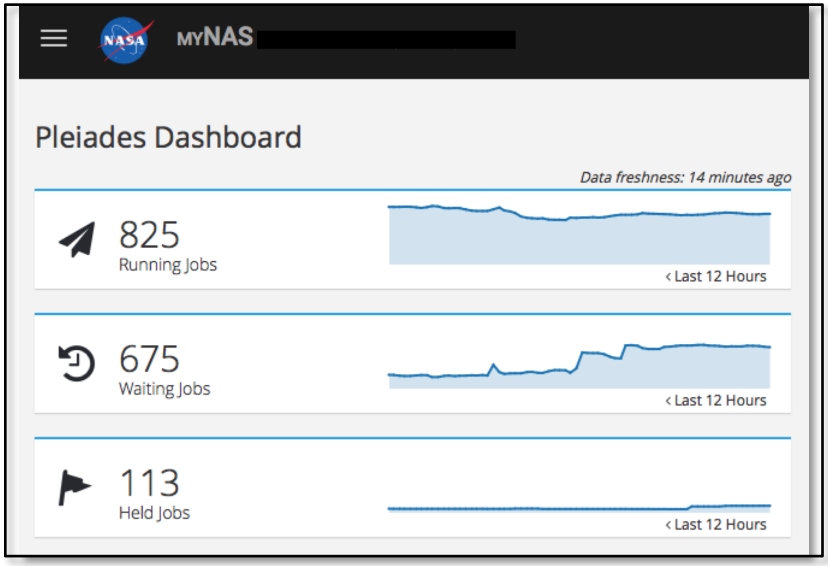

Numerous charts on the myNAS website allow users to obtain an overall summary of jobs or to drill down for detailed information.

POC: John Hardman, john.hardman@nasa.gov, (650) 604-0417, NASA Advanced Supercomputing Division, CSRA LLC

New Tool Developed for Analyzing MPI and I/O Performance

- HECC staff developed a performance profiling tool, MPIProf, to analyze the MPI and I/O performance of applications running on HECC systems. The tool reports aggregated and per-rank profiling results:
  - Function timing, number of calls, message size
  - MPI I/O and POSIX I/O statistics
  - Call-path based information
  - Memory usage
- Compared to other tools, MPIProf has a simple interface and low overhead, and works with many MPI implementations. The in-house development of the tool allows quick testing of new capabilities.
- On August 24, a NAS training webinar was conducted on the use of MPIProf, and was well received by users.
- Staff members are successfully using the tool in conjunction with other tools to analyze the top 10 applications running on Pleiades (see slide 6). This effort has resulted in identification of performance bottlenecks in a few applications, which enabled subsequent optimization for substantial performance improvements.

Mission Impact: Understanding performance characteristics of applications is very important for reducing time-to-solution and improving HECC compute resource utilization. Effective performance analysis tools are a critical component of the process.

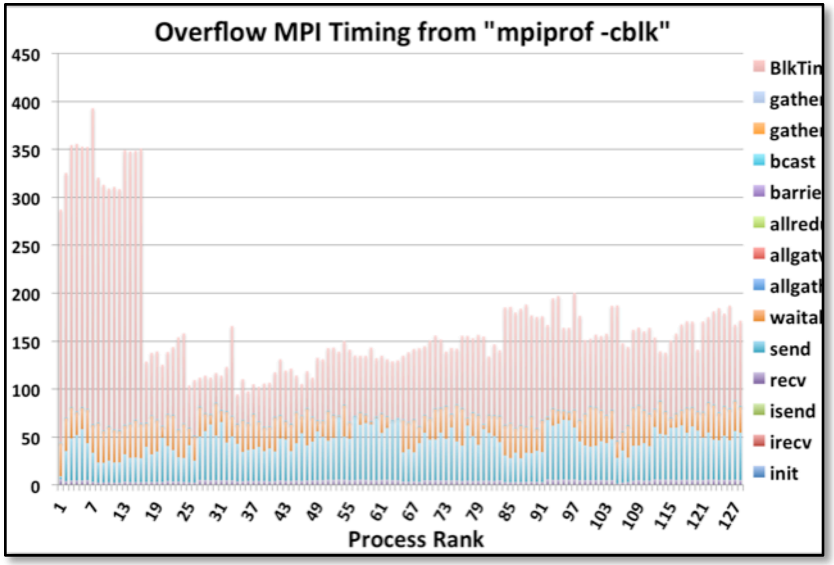

A sample time profiling of MPI functions called in the OVERFLOW code for each MPI rank, as reported by MPIProf. Different color bars indicate different MPI functions. This test case was run with 128 processes on 8 Sandy Bridge nodes of Pleiades. The graph shows load imbalance in the first 16 MPI ranks, which impacts overall performance of the application.

POC: Henry Jin, haoqiang.jin@nasa.gov, (650) 604-0165, NASA Advanced Supercomputing Division

I/O Characterization of NASA Applications Identifies Scaling Improvements

- Using the in-house MPIProf tool (see slide 5) and the IOT toolkit, HECC Application Performance and Productivity (APP) staff characterized the I/O of 14 applications running on Pleiades, 8 of which are among the top 10 on Pleiades in terms of Standard Billing Units (SBU)* usage.
- MPIProf is used for MPI and I/O profiling, while IOT (licensed from IO Doctors, LLC) is specialized for I/O analysis and optimization.
- Analysis results show that many of the 14 applications exhibit common I/O characteristics:
  - Rank 0 dominates the I/O.
  - I/O sizes less than or equal to 4 megabytes.
  - Contain more writes than reads.
  - Do sequential I/O operations.
- For an ATHENA-C++ application that uses MPI-IO, APP staff identified the I/O bottleneck and sped up I/O by 10x for a 3,264-rank case and 36x for a 5,968-rank case. The speedup will help improve the scaling of this application to larger processor counts.

** 1 SBU equals 1 hour of a Pleiades Westmere 12-core node.*

Mission Impact: I/O characterization establishes the I/O landscape of applications and helps HECC plan for future I/O infrastructure. It also helps HECC staff in identifying potential bottlenecks for improvement in applications.

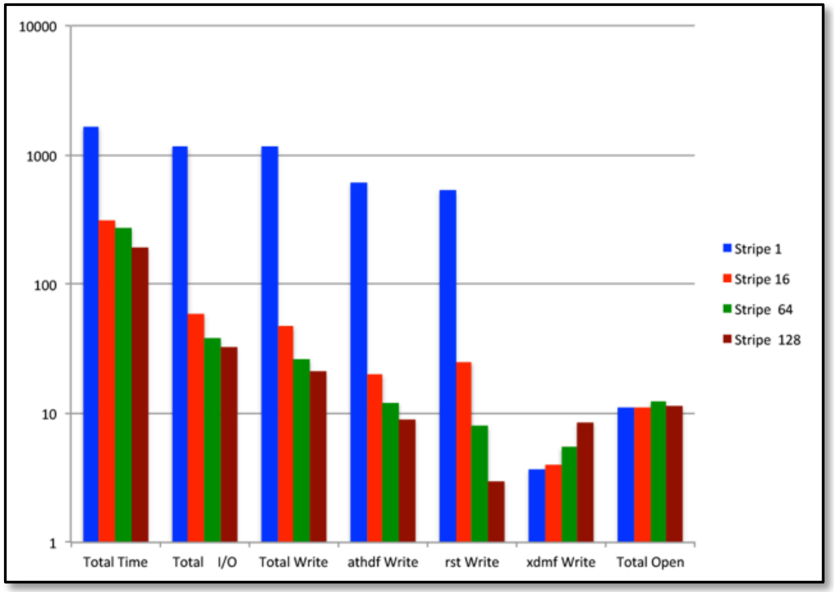

Chart showing timing for an ATHENA-C++ application (5,968-rank case) for the whole run, I/O open and write time for the whole run, and write time for certain files (athdf, rst, xdmf) with different stripe counts. (**View in Slide Show mode and click image for full scale.**)

POCs: Sherry Chang, sherry.chang@nasa.gov, (650) 604-1272, NASA Advanced Supercomputing (NAS) Division, CSRA LLC; Henry Jin, haoqiang.jin@nasa.gov, (650) 604-0165, NAS Division

Linked-Scatter-Plot System on hyperwall Used to Examine SETI Institute Data

- HECC's Visualization team significantly enhanced their existing linked-scatter-plot (LSP) framework to help the SETI Institute analyze radio-wave observations in its database.
- LSP enables insightful views into large datasets, which, in this case, allows rapid and accurate searches for extraterrestrial signals in our galaxy. The extensible core is customized for the unique challenges of this particular dataset.
  - SETI has gathered large numbers of observations, but had not previously done deep analysis and postprocessing of that data. To date, they have had three half-day working sessions with LSP on the hyperwall to classify data.
  - The Visualization team wrote three custom viewers for SETI "waterfall plot" data, and integrated five signal-processing methods and up to 2,000 CPU cores into the framework, allowing interactive processing of user-specified data subsets.
- Based on the Visualization team's enhancements, the SETI team found several signals of interest and gained greater insight into their data.

Mission Impact: The LSP system gives SETI researchers a powerful tool to examine, classify, and process, large datasets to gain insight and understanding into their data.

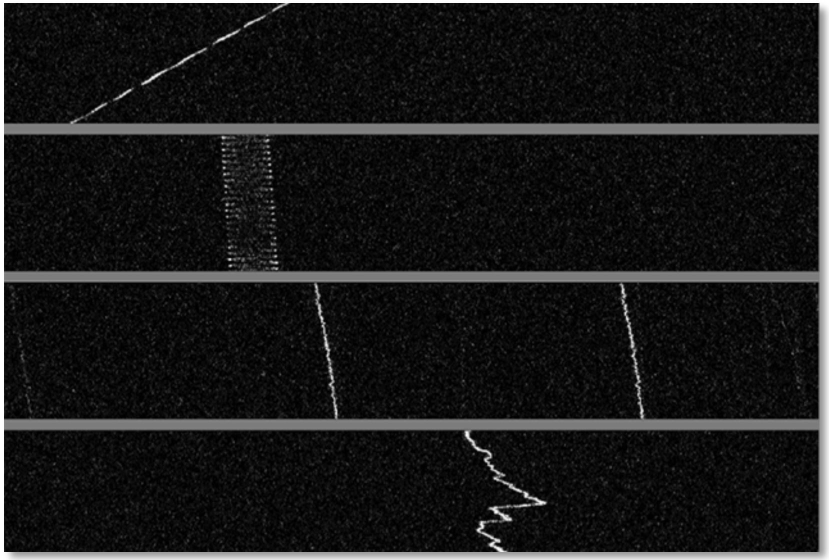

This image produced by HECC visualization experts shows four examples of time-varying spectrograms, also called "waterfall plots," from the SETI dataset. The full dataset is approximately 335,000 primary observations; each primary observation consists of 16 sub-bands, for a total of over 5 million waterfall plots.

POCs: Chris Henze, chris.henze@nasa.gov (650) 604-3959, NASA Advanced Supercomputing (NAS) Division; Bron Nelson, bron.c.nelson@nasa.gov, (650)604-4329, NAS Division, CSRA LLC

Facilities Team Successfully Coordinates Demolition of Old N258 Cooling Tower

- HECC Facilities engineers, in coordination with NASA Ames' Code J facilities engineers, completed the demolition of building N258's original, 30-year-old cooling tower.
- A replacement cooling tower was completed and put into production in May, and the old tower remained in place for a period of time in case a failback system was needed.
- Through detailed planning and coordination, the Facilities team ensured that all safety, environmental, and engineering challenges were addressed without incident and without affecting HECC supercomputing production.
- The new cooling tower provides added redundancy (four independent cells, each of which can be taken offline without affecting production) and 50% more cooling capacity.

Mission Impact: The HECC Facilities team's careful planning and coordination with other organizations ensured that the old cooling tower was demolished with no impact on safety, environmental requirements, or HECC user productivity.

Photo of the demolition of the 30-year-old cooling tower associated with the NASA Advanced Supercomputing facility at NASA Ames. The tower was replaced with a modern cooling system in May 2016.

POCs: Chris Buchanan, chris.buchanan@nasa.gov, (650) 604-4308,
Chris Tanner, christopher.tanner@nasa.gov, (650) 604-6754,
NASA Advanced Supercomputing (NAS) Division, CSRA LLC

Space Launch System Booster Separation Simulations Run on Pleiades

- The aerodynamic forces acting on the two SLS solid rocket boosters (SRBs) play an important role in determining the flight of the boosters as they separate from the center core.
- Aerospace engineers at NASA Ames built the most complex computational fluid dynamics (CFD)-based aerodynamic database in the SLS program to date, to help verify the SLS core and booster design, utilizing CFD numerical simulations.
- The team used CFD data, along with results from a NASA Langley wind tunnel test, to run simulations of 22 plumes, including 16 booster-separation motors (BSMs):
  - Ran more than 22K Cart3D simulations and 390 OVERFLOW simulations in two months.
  - Built an aerodynamic database to cover 8 independent variables, including booster translation and rotation variables, free-stream angle of attack, angle of side-slip, and relative thrust of BSMs.
- All simulations were run on the Pleiades supercomputer, primarily on the 24-core Intel Xeon “Haswell” nodes.

Mission Impact: The Guidance, Navigation, and Control group at NASA Marshall will use this database to run simulations to ensure successful SLS booster separation under all possible flight conditions.

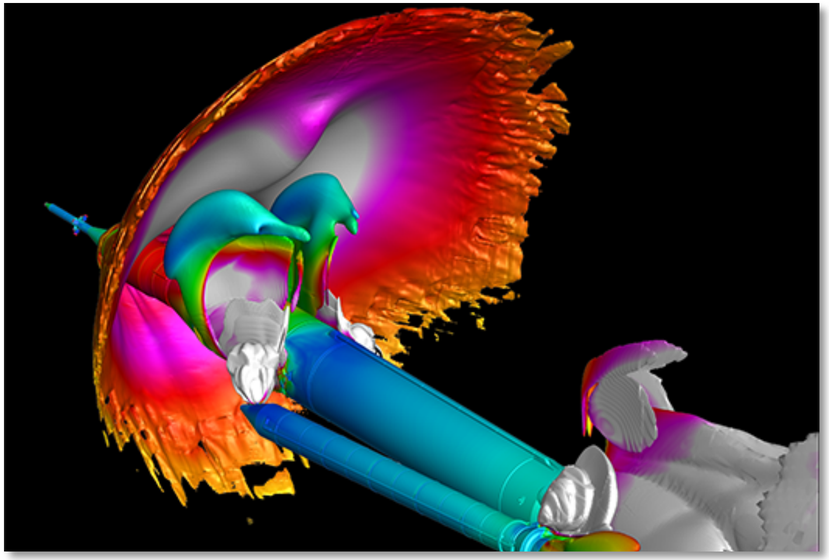

SLS booster separation flow field: isosurfaces of density illustrate a shock wave on the front of the vehicle; isosurfaces of plume-gas species show the booster separation-motor plumes. *Stuart Rogers, NASA/Ames; Ryan Rocha, University of California, Davis*

POCs: Stuart Rogers, stuart.e.rogers@nasa.gov, (650) 604-4481, NASA Advanced Supercomputing (NAS) Division; Derek Dalle, derek.j.dalle@nasa.gov, (650) 604-4209, NAS Division, Science and Technology Corp.

** HECC provided supercomputing resources and services in support of this work*

Pleiades Enables Sensitivity Analysis for Chaotic Fluid Simulations

- The FUN3D development team at NASA Langley, with researchers at MIT, ran compute-intensive simulations on Pleiades to evaluate and implement their new Least Squares Shadowing (LSS) approach for performing sensitivity analysis of chaotic fluid behavior.
- LSS solves a globally coupled version of linearized equations to determine a bounded adjoint solution that can be used to compute sensitivity information for aircraft design optimization, error estimation, and uncertainty quantification.
  - As a baseline test, the team simulated shedding flow over a 2D NACA 0012 airfoil at high angle of attack, with chaotic shedding from the wing's upper surface.
  - The team successfully computed—for the first time—a bounded discrete adjoint solution for this flow using the LSS method, overcoming a traditional solution problem where discrete adjoint equations grew exponentially over time.
- The team is working to extend the LSS approach to large-scale 3D simulations, which would have a broad impact on fluid simulations of relevance to NASA aerospace work across the speed range.

** HECC provided supercomputing resources and services in support of this work*

Mission Impact: Researchers required a large portion of HECC's Pleiades supercomputer in order to perform Least Squares Shadowing (LSS) simulations to obtain bounded adjoint solutions for a trivial CFD analysis problem that fits easily on a single CPU core. NASA will require many orders of magnitude more compute power than currently available to apply the LSS method to mission-relevant problems.

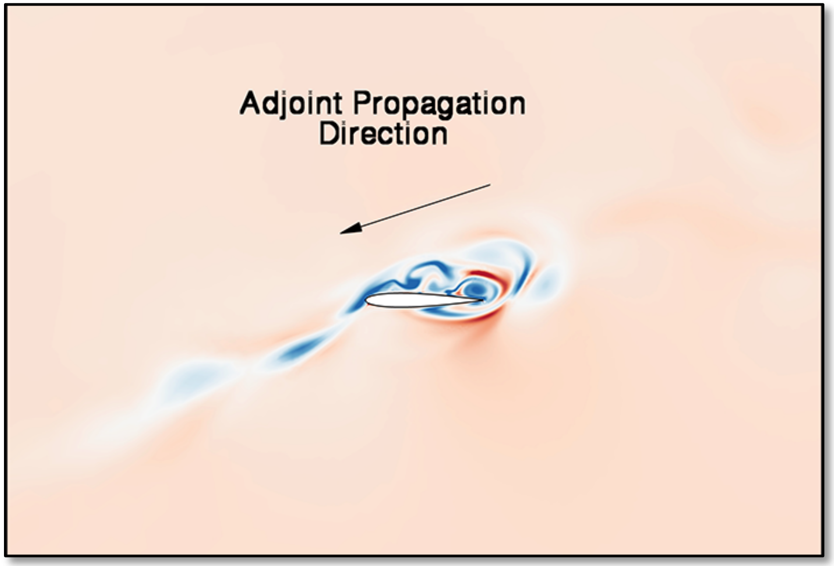

Snapshot of a discrete adjoint solution computed using the Least Squares Shadowing (LSS) technique.

POC: Eric Nielsen, eric.j.nielsen@nasa.gov, (757) 864-2239, NASA Langley Research Center

HECC Facility Hosts Several Visitors and Tours in August 2016

- HECC hosted 19 tour groups in August; guests learned about the agency-wide missions being supported by HECC assets, and some groups also viewed the D-Wave 2X quantum computer system. Tours were limited this month due to a facility shutdown. Visitors this month included:
  - Jack Kaye, Associate Director for Research of the Earth Science Division within the Science Mission Directorate.
  - Jeannette Ruiz-Hanna, NASA's Associate Chief Information Officer for IT Security, and Senior Agency Information Security Official, (NASA Headquarters).
  - Sean Gallagher, Chief Information Officer, (NASA Glenn), and Rob Powell, Senior adviser for Cybersecurity (NASA Headquarters).
  - Vernon Gibson, Chief Scientific Advisor to the United Kingdom Ministry of Defense.
  - Sam Black, Program Examiner at the White House Office of Management and Budget.
  - Staff from the Royal Institute of Technology (KTH) Innovation Office.
  - LaNetra Tate, Program Executive, Game Changing Development Program, NASA Headquarters.
  - Attendees of the Thermal and Fluids Analysis Workshop.
  - A group of local students from Gavilan Community College.

NAS Division aerospace engineer Shishir Pandya presented computational results to attendees of the 2016 Thermal and Fluids Analysis Workshop hosted at NASA Ames; Pandya was conference chair of this year's meeting.

POC: Gina Morello, gina.f.morello@nasa.gov, (650) 604-4462, NASA Advanced Supercomputing Division

Papers

- "**The IBEX Ribbon and the Pickup Ion Ring Stability in the Outer Heliosheath. I. Theory and Hybrid Simulations**," V. Florinski, J. Heerikhuisen, J. Niemiec, A. Ernst, The Astrophysical Journal, Vol. 826, No. 2, July 29, 2016. *
  http://iopscience.iop.org/article/10.3847/0004-637X/826/2/197/meta
- "**The IBEX Ribbon and the Pickup Ion Ring Stability in the Outer Heliosheath. II. Monte-Carlo and Particle-in-Cell Model Results**," J. Niemiec, V. Florinski, J. Heerikhuisen, K.-I. Nishikawa, The Astrophysical Journal, Vol. 826, No. 2, July 29, 2016. *
  http://iopscience.iop.org/article/10.3847/0004-637X/826/2/198/meta
- "**Capture of Trans-Neptunian Planetesimals in the Main Asteroid Belt**," D. Vokrouhlický, W. F. Bottke, D. Nesvorný, The Astronomical Journal, Vol. 152, No. 2, August 2016. *
  http://iopscience.iop.org/article/10.3847/0004-6256/152/2/39/meta
- "**Optimizing the Energy Consumption of Message Passing Applications with Iterations Executed Over Grids**," A. Fanfakh, J.-C. Charr, R. Courtuier, A. Giersch, Journal of Computational Science, August 7, 2016. *
  http://www.sciencedirect.com/science/article/pii/S187775031630120X
- "**On the Electron Agyrotropy During Rapid Asymmetric Magnetic Island Coalescence in Presence of a Guide Field**," E. Cazzola, M. Innocenti, M. Goldman, D. Newman, S. Markdis, G. Lapenta, Geophys. Res. Letters, Vol. 43, Issue 15, August 16, 2016. *
  http://onlinelibrary.wiley.com/doi/10.1002/2016GL070195/full

** HECC provided supercomputing resources and services in support of this work*

Papers (cont.)

- "**Grounding Line Variability and Subglacial Lake Drainage on Pine Island Glacier, Antarctica**," I. Joughin, D. Shean, B. Smith, P. Dutrieux, Geophysical Research Letters, August 17, 2016. *
  http://onlinelibrary.wiley.com/doi/10.1002/2016GL070259/full
- "**Helioseismic Holography of Simulated Sunspots: Magnetic and Thermal Contributions to Travel Times**," T. Felipe, D. Braun, A. Crouch, A. Birch, arXiv:1608.04893 [astro-ph.SR], August 17, 2016. *
  http://arxiv.org/abs/1608.04893
- "**Application of the Metabolic Scaling Theory and Water-Energy Balance Equation to Model Large-Scale Patterns of Maximum Forest Canopy Height**," S. Choi, S. Ganguly, et al., Global Ecology and Biogeography, August 18, 2016. *
  http://onlinelibrary.wiley.com/doi/10.1111/geb.12503/full
- "**Evaluating Summer-Time Ozone Enhancement Events in the Southeast United States**," M. Johnson, S. Kuang, L. Wang, M. Newchurch, Atmosphere, Vol. 7, August 19, 2016. *
  http://www.mdpi.com/2073-4433/7/8/108/htm
- "**Mesoscale Modulation of Air-Sea CO2 Flux in Drake Passage**," H. Song, et al., Journal of Geophysical Research: Oceans, August 24, 2016. *
  http://onlinelibrary.wiley.com/doi/10.1002/2016JC011714/full

** HECC provided supercomputing resources and services in support of this work*

News and Events

- **Endeavour Simulations Help Reshape the Future of Aircraft Fuel Efficiency**, *NAS Feature Story*, August 19, 2016—Flexible fixed wing simulations performed on HECC's Endeavour supercomputer at the NASA Advanced Supercomputing Division give a deeper look into how flexible wing shapes can increase fuel efficiency by adapting to varying flight conditions.
  http://www.nas.nasa.gov/publications/articles/feature_Flaps_VCCTEF_Aftosmis.html
- **NAS-Developed Pegasus 5 Wings Its Way to NASA's Software of the Year**, *NAS Feature Story*, August 26, 2016—NASA awarded the Pegasus 5 CFD tool, developed by a team led by NAS aerospace engineer Stuart Rogers, with the 2016 Software of the Year Award.
  http://www.nas.nasa.gov/publications/articles/feature_Pegasus_Software_of_the_Year.html

HECC Utilization

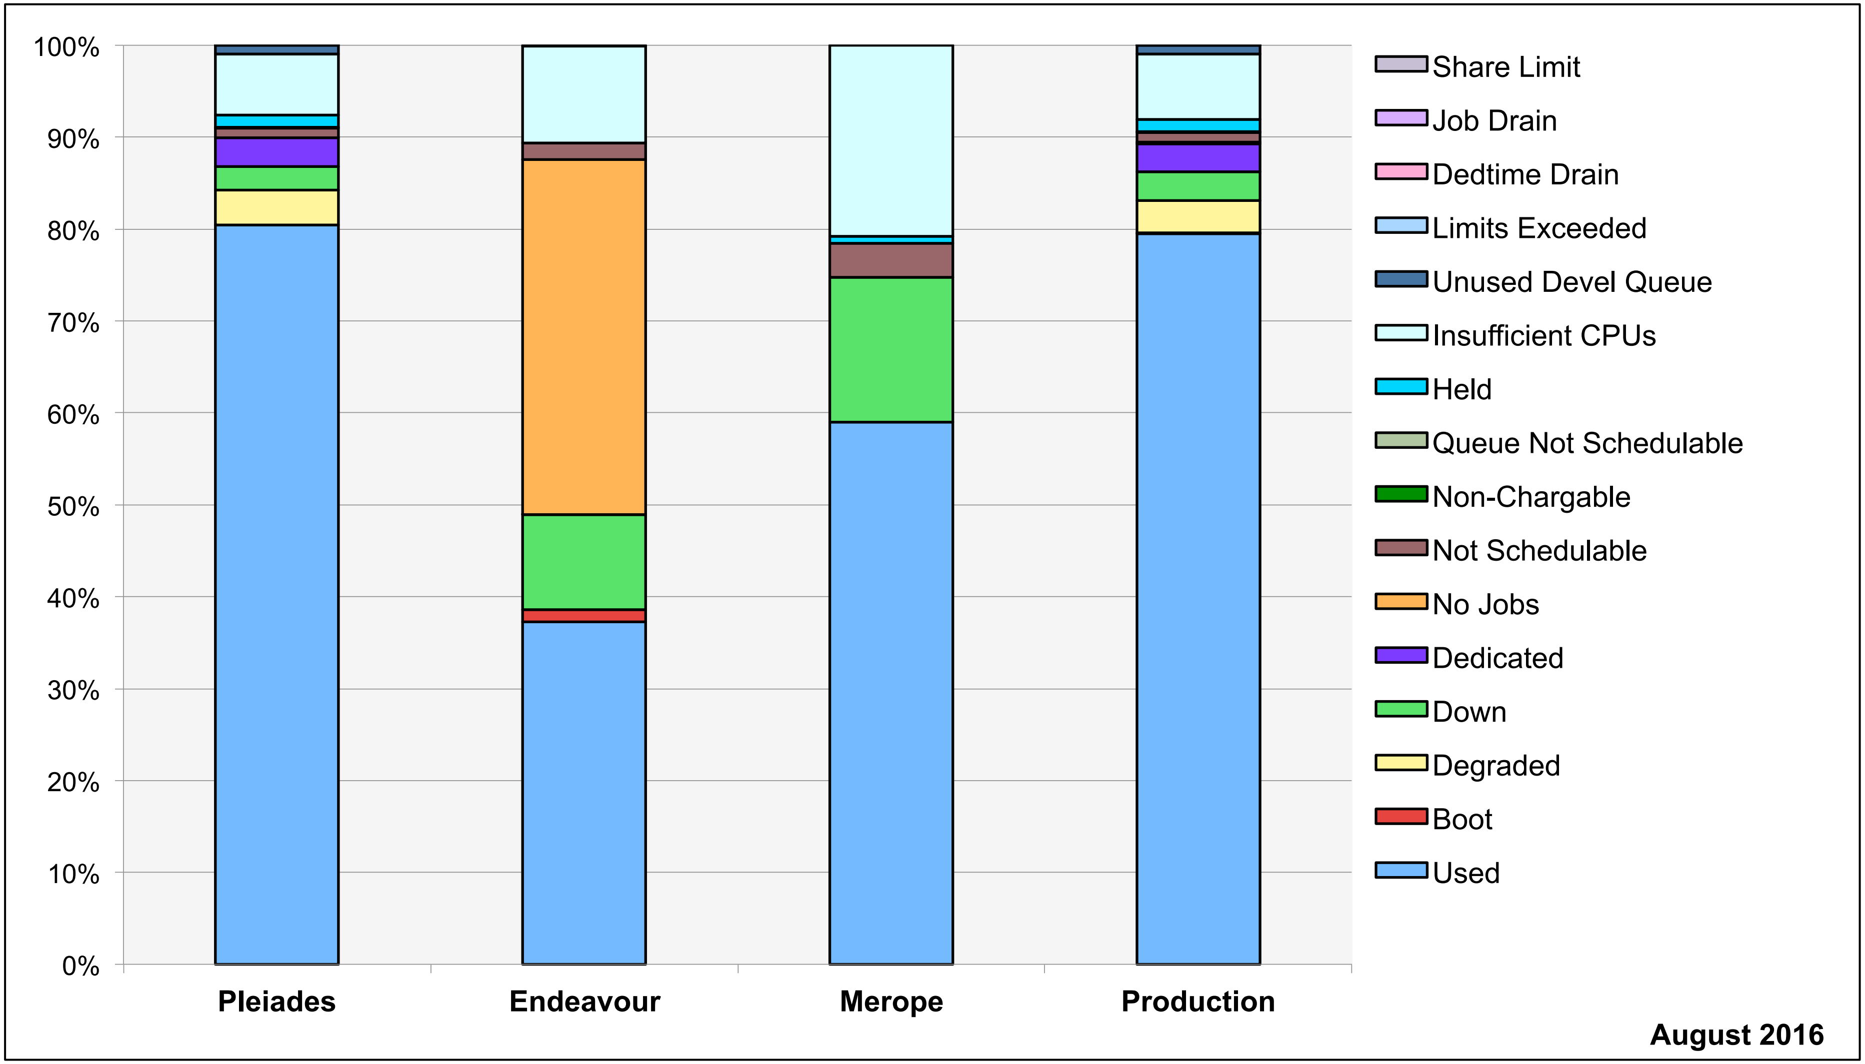

HECC Utilization Normalized to 30-Day Month

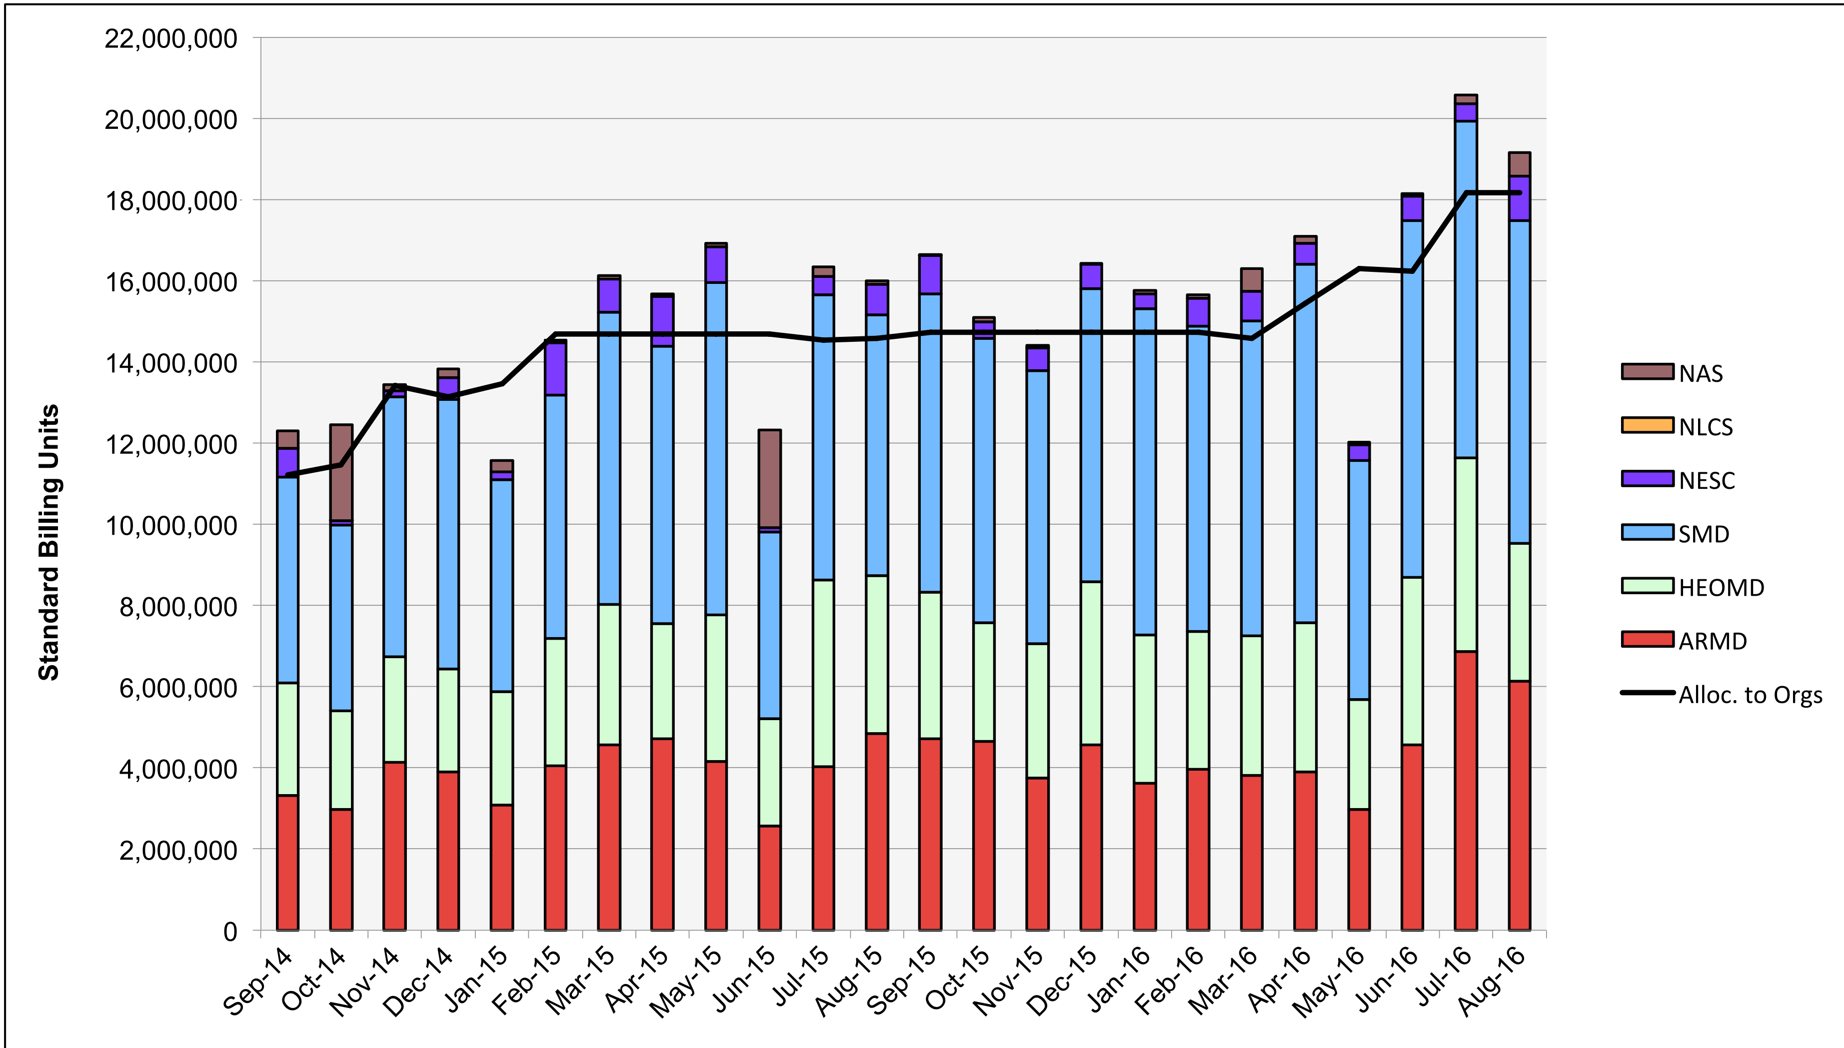

HECC Utilization Normalized to 30-Day Month

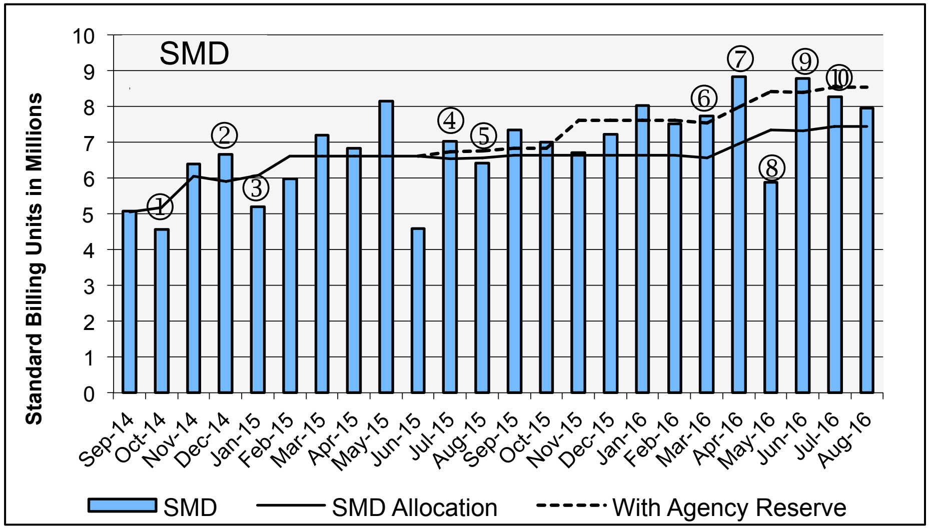

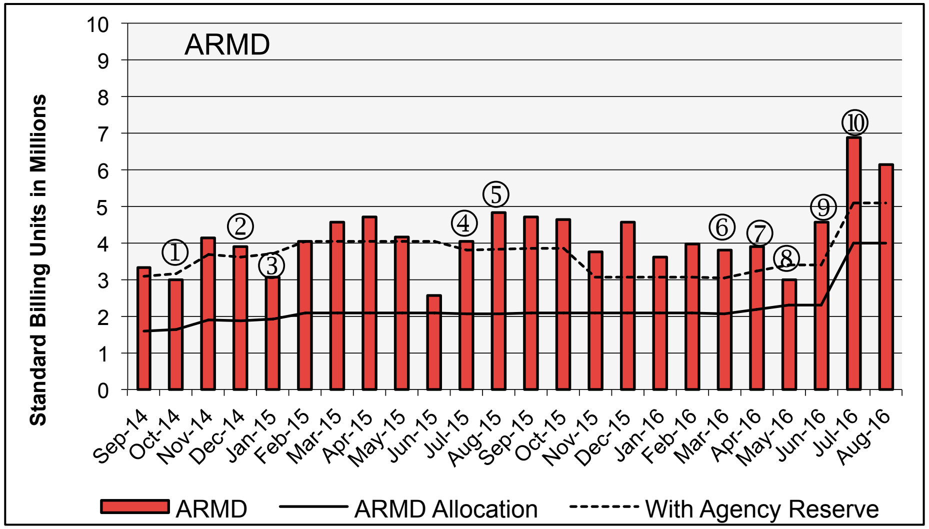

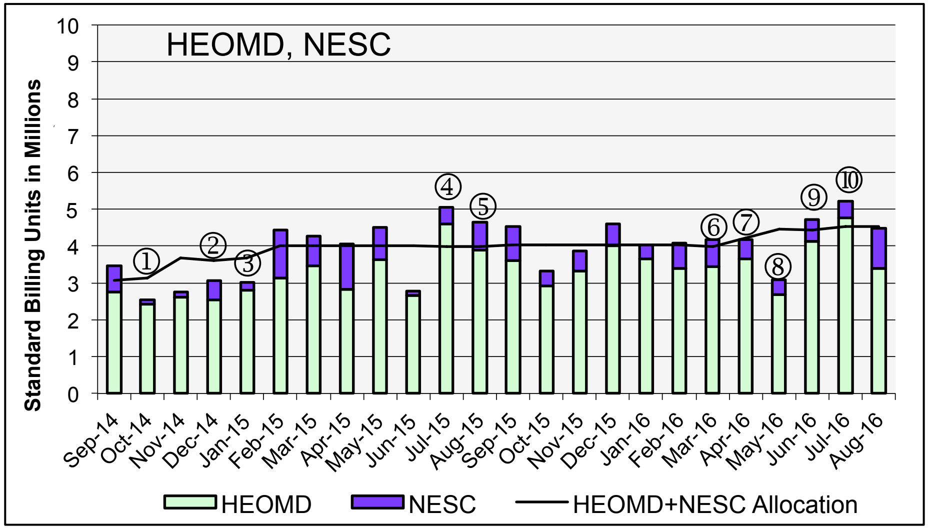

① 16 Westmere racks retired from, 3 Ivy Bridge and 15 Haswell racks added to Pleiades; 10 Nehalem and 2 Westmere racks added to Merope
② 16 Westmere racks retired from Pleiades
③ 14 Haswell racks added to Pleiades
④ 7 Nehalem ½ racks retired from Merope
⑤ 7 Westmere ½ racks added to Pleiades
⑥ 16 Westmere racks retired from Pleiades
⑦ 10 Broadwell racks added to Pleiades
⑧ 4 Broadwell racks added to Pleiades
⑨ 14 (all) Westmere racks retired from Pleiades
⑩ 14 Broadwell racks added to Pleiades (10.5 Dedicated to ARMD)

Tape Archive Status

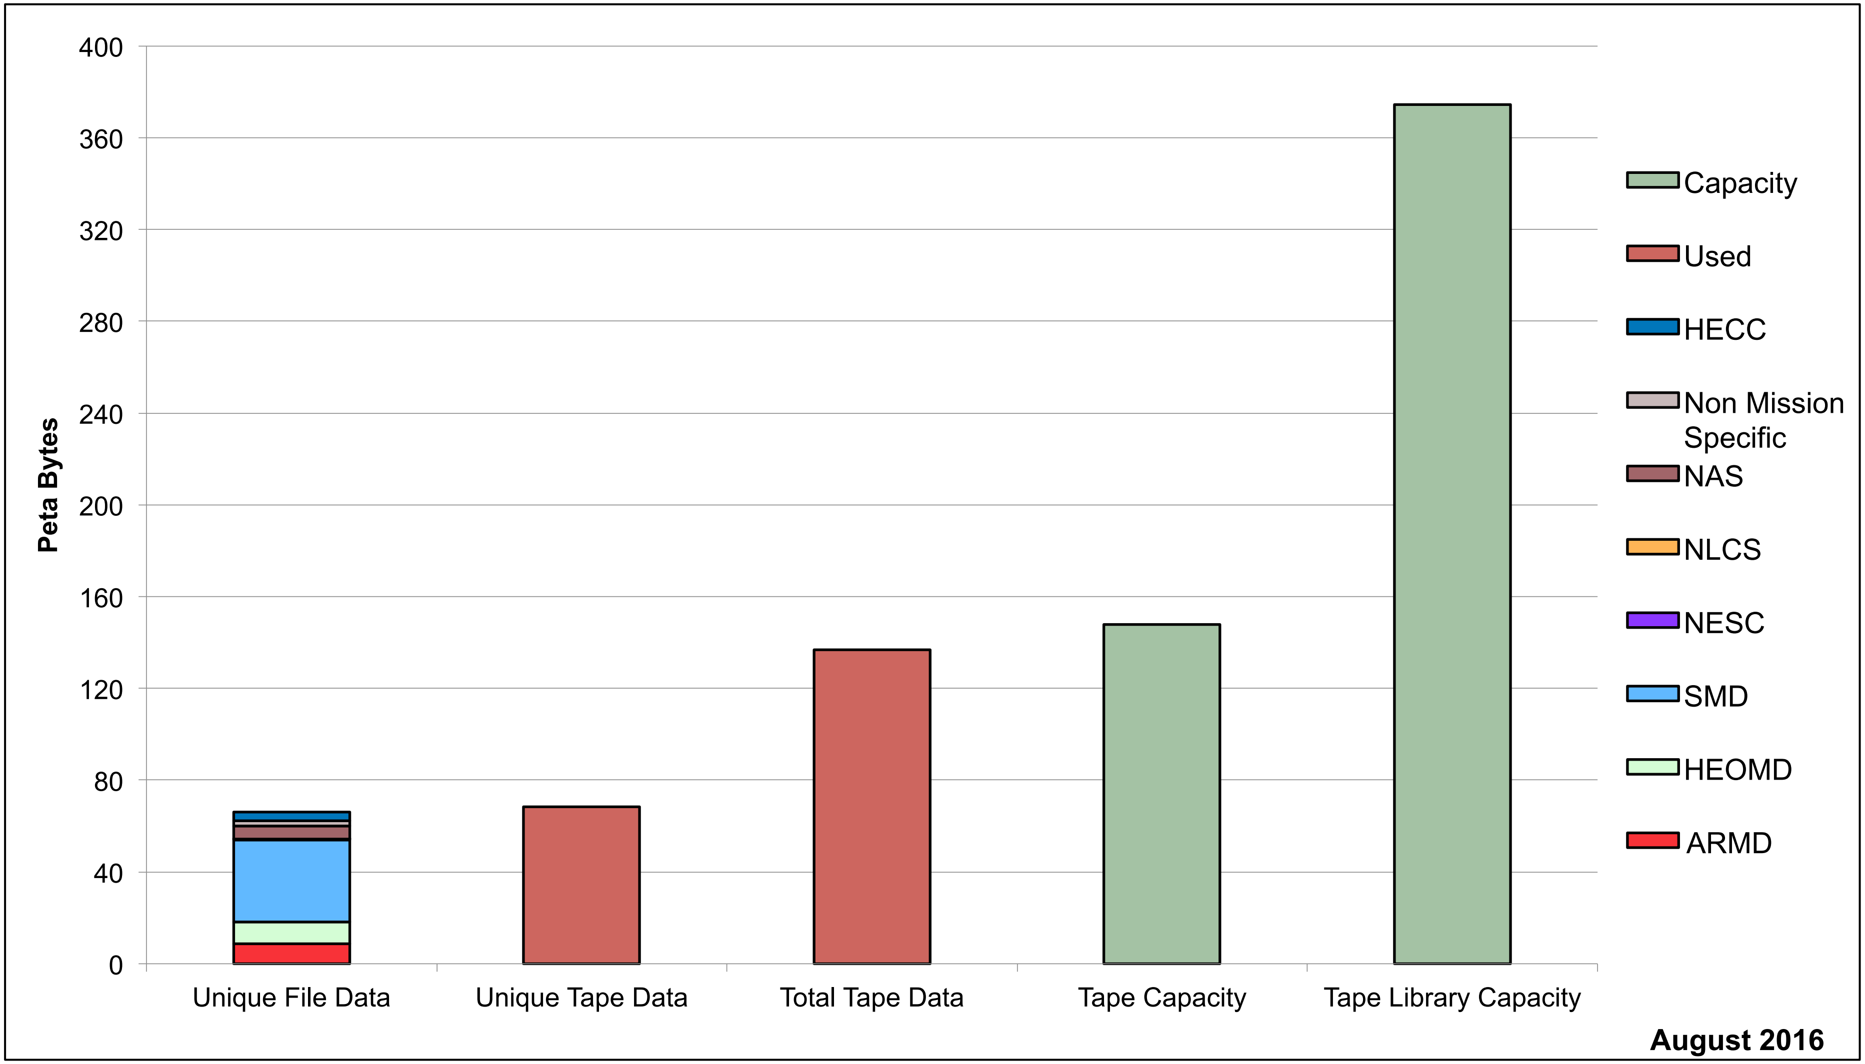

Tape Archive Status

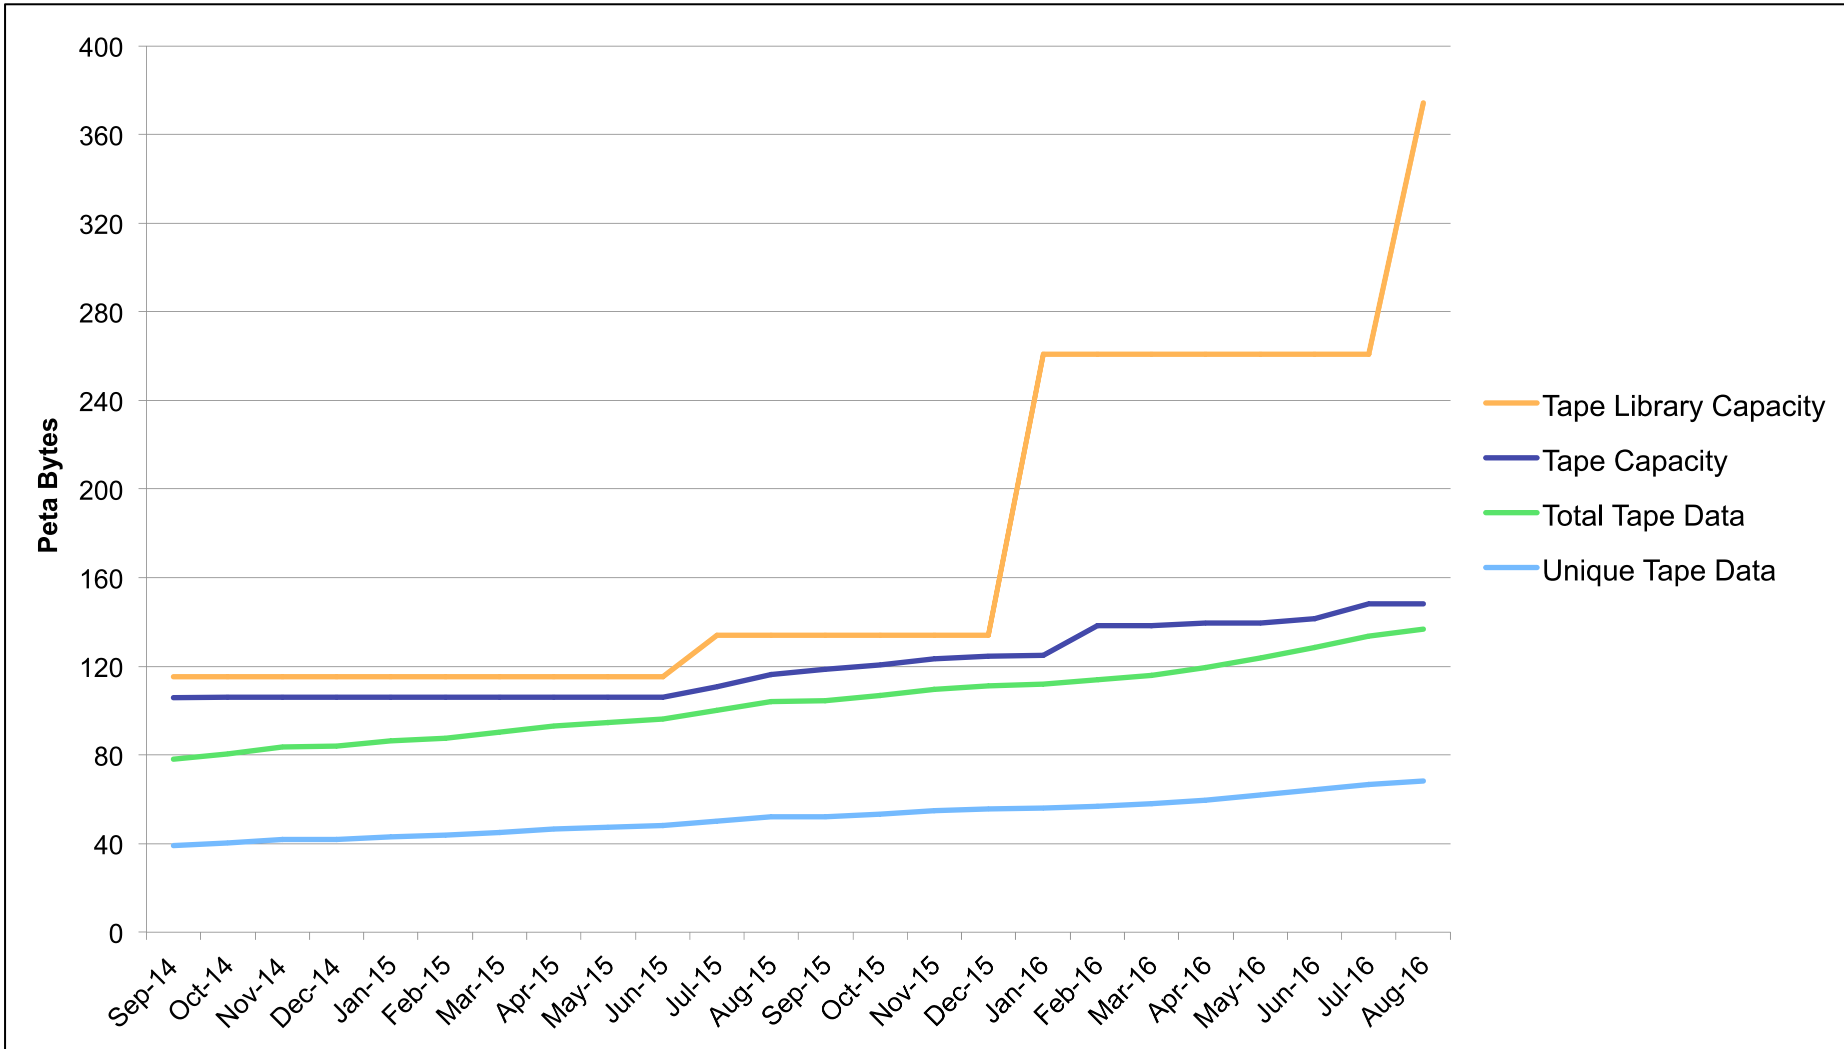

Pleiades:
SBUs Reported, Normalized to 30-Day Month

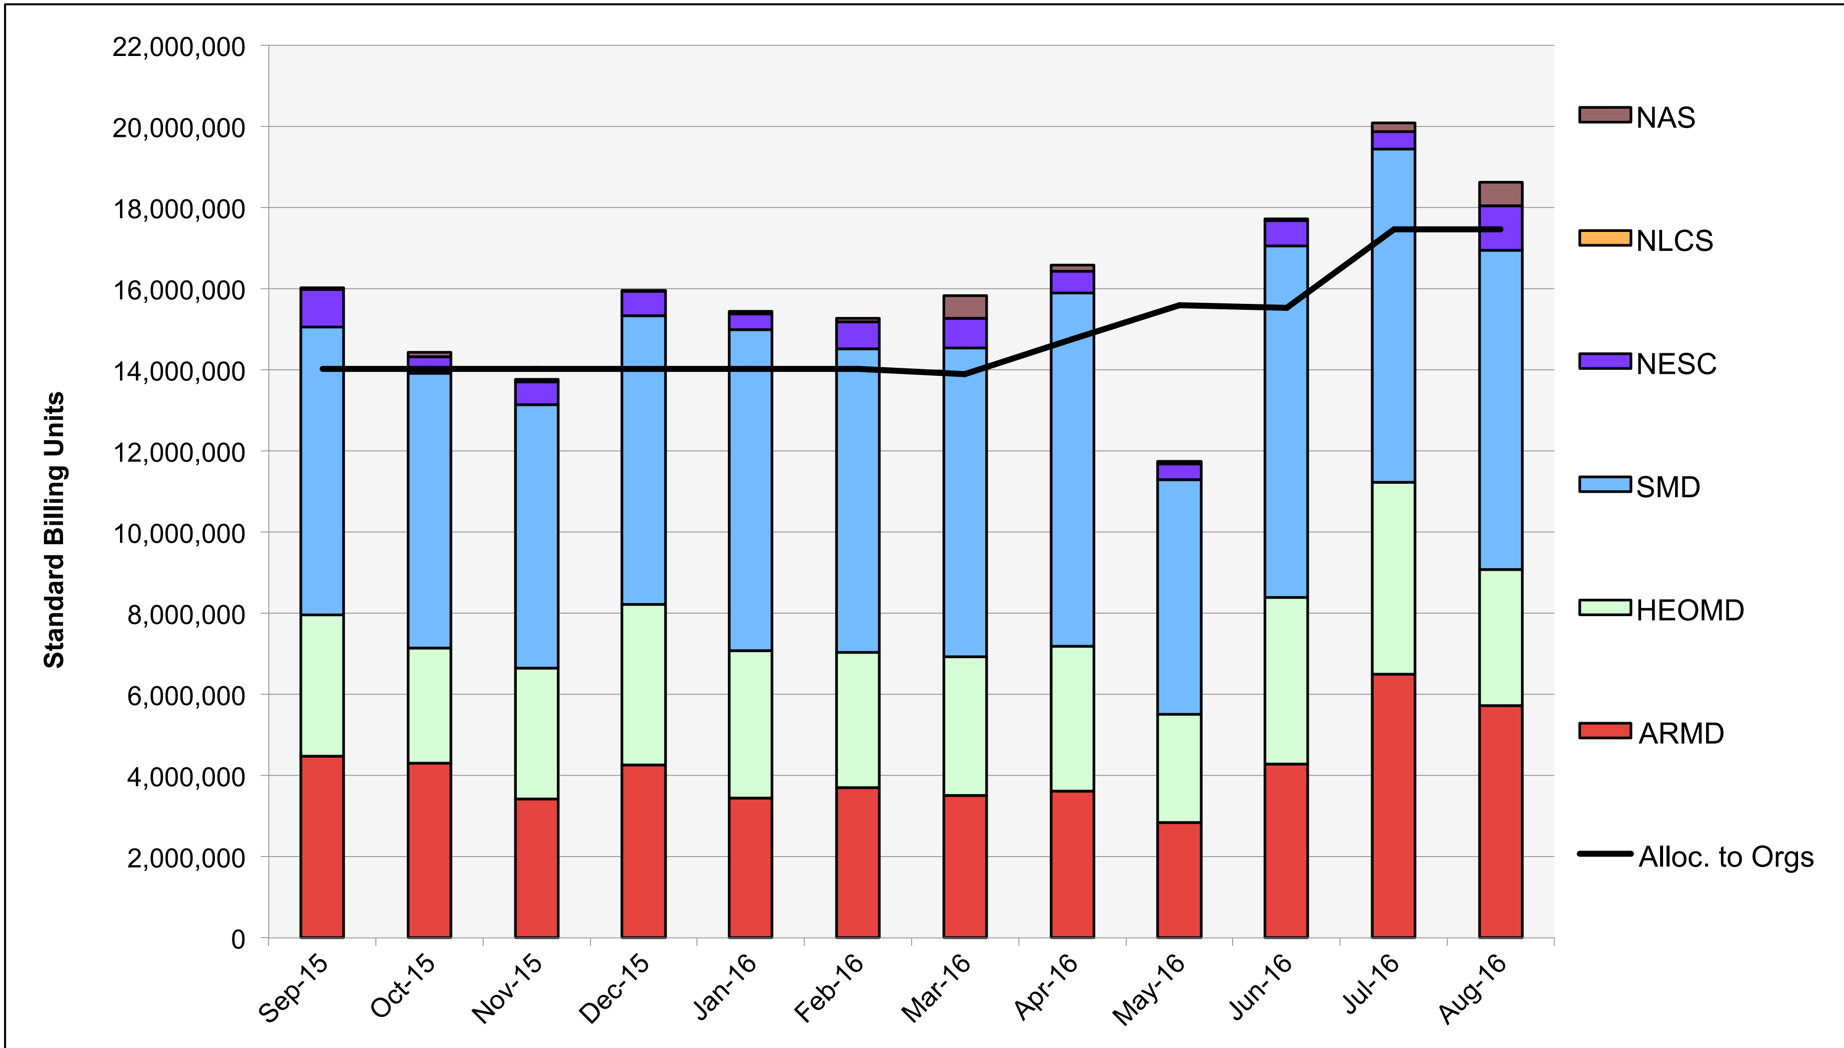

Pleiades: Devel Queue Utilization

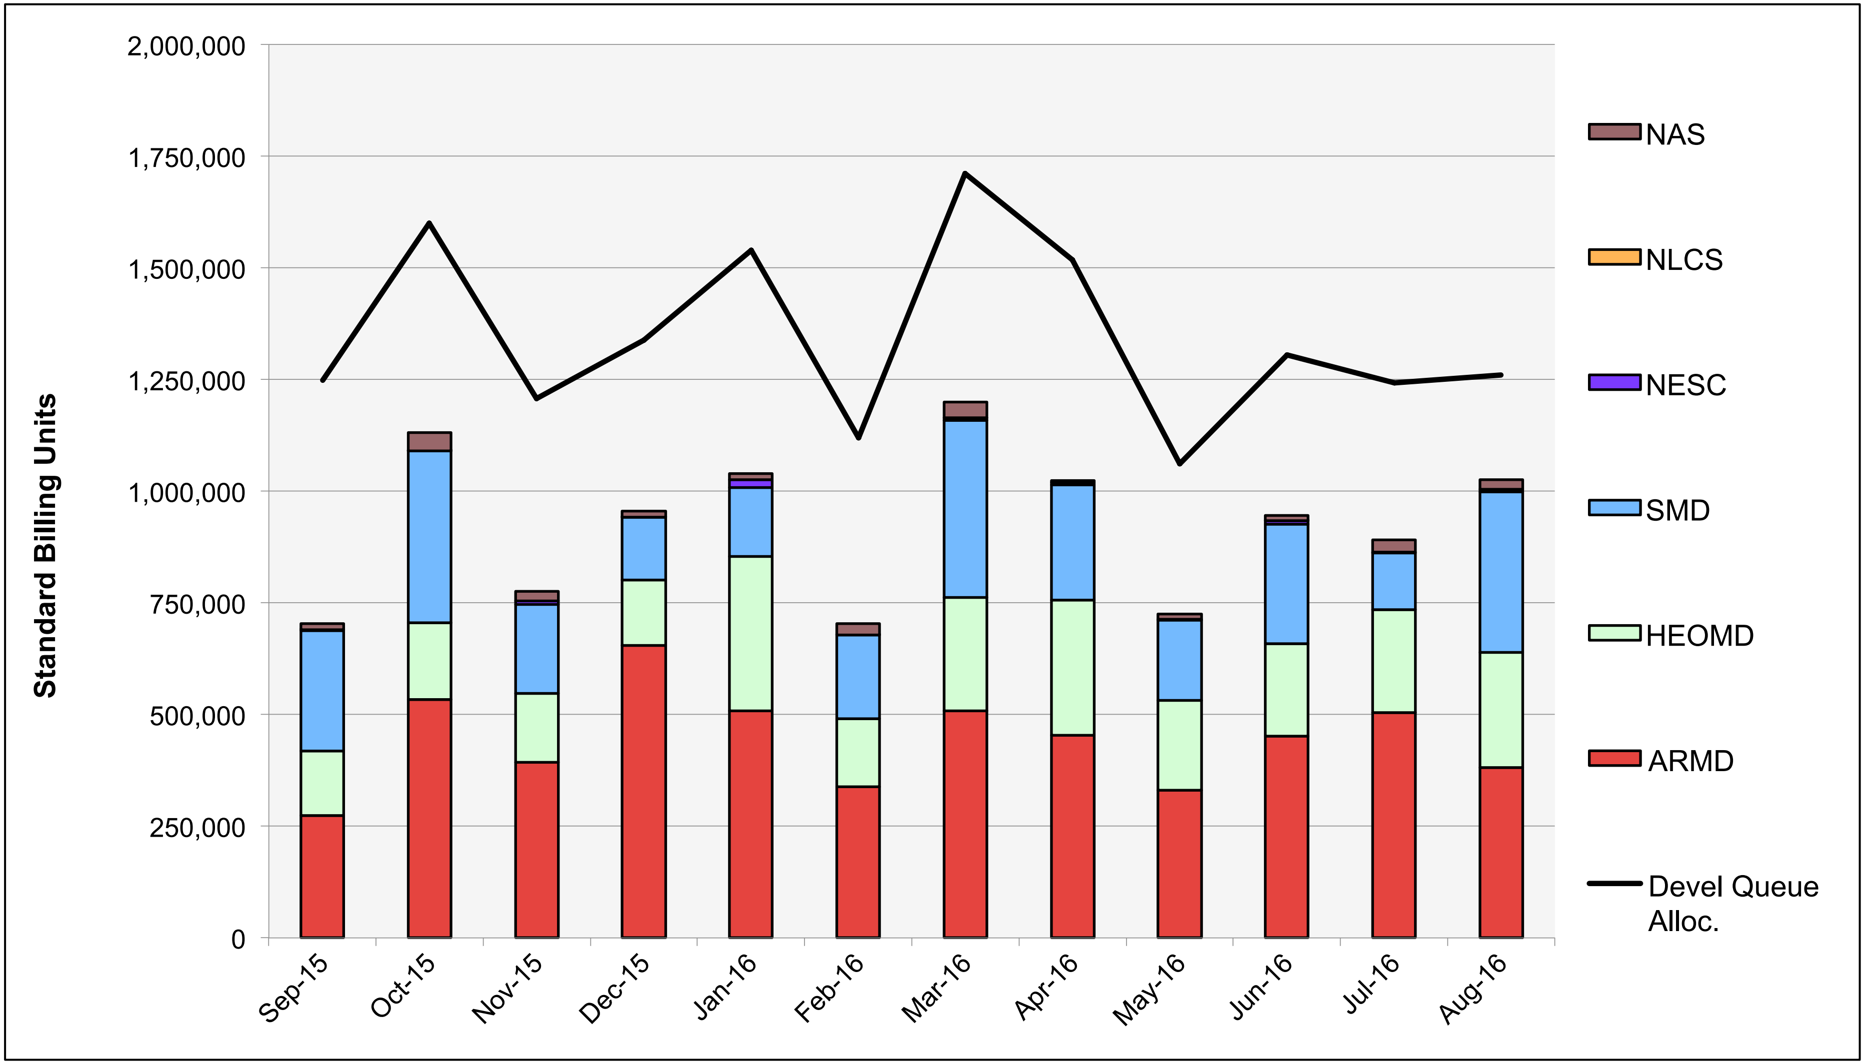

Pleiades: Monthly Utilization by Job Length

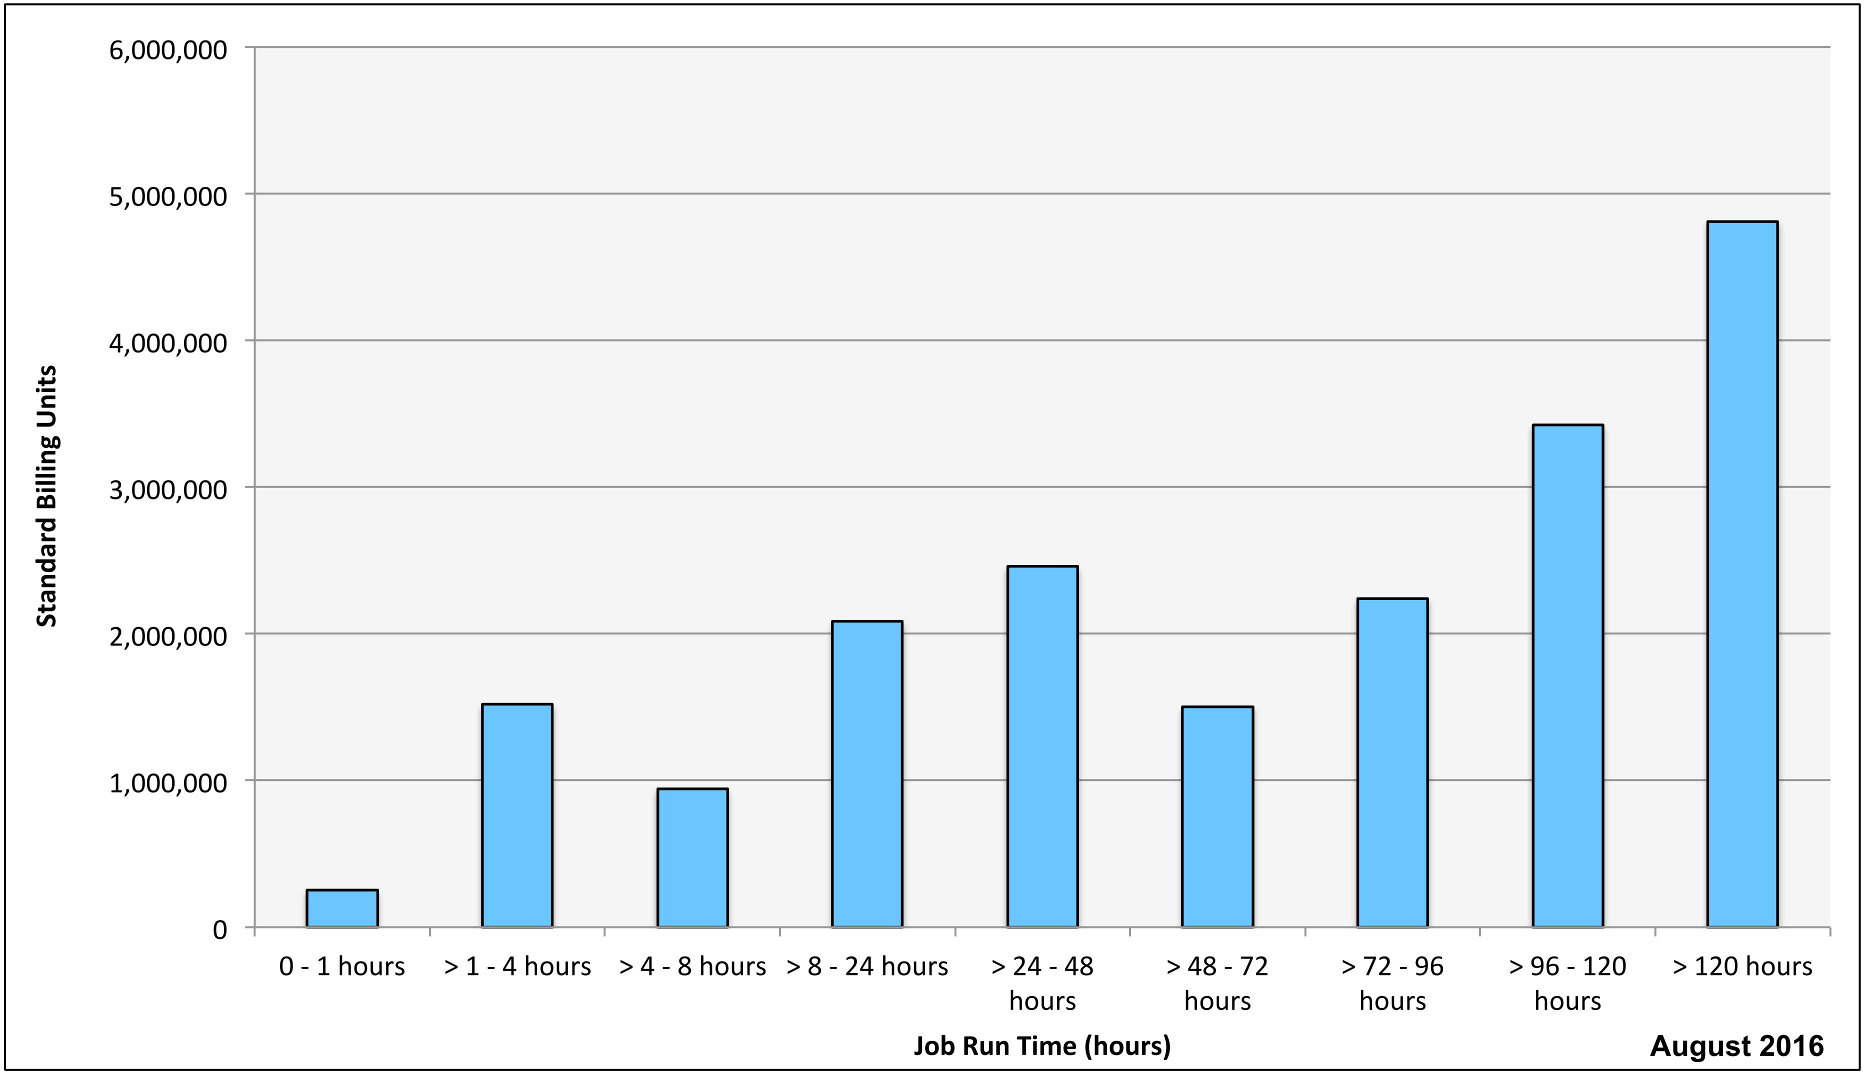

Pleiades:
Monthly Utilization by Size and Mission

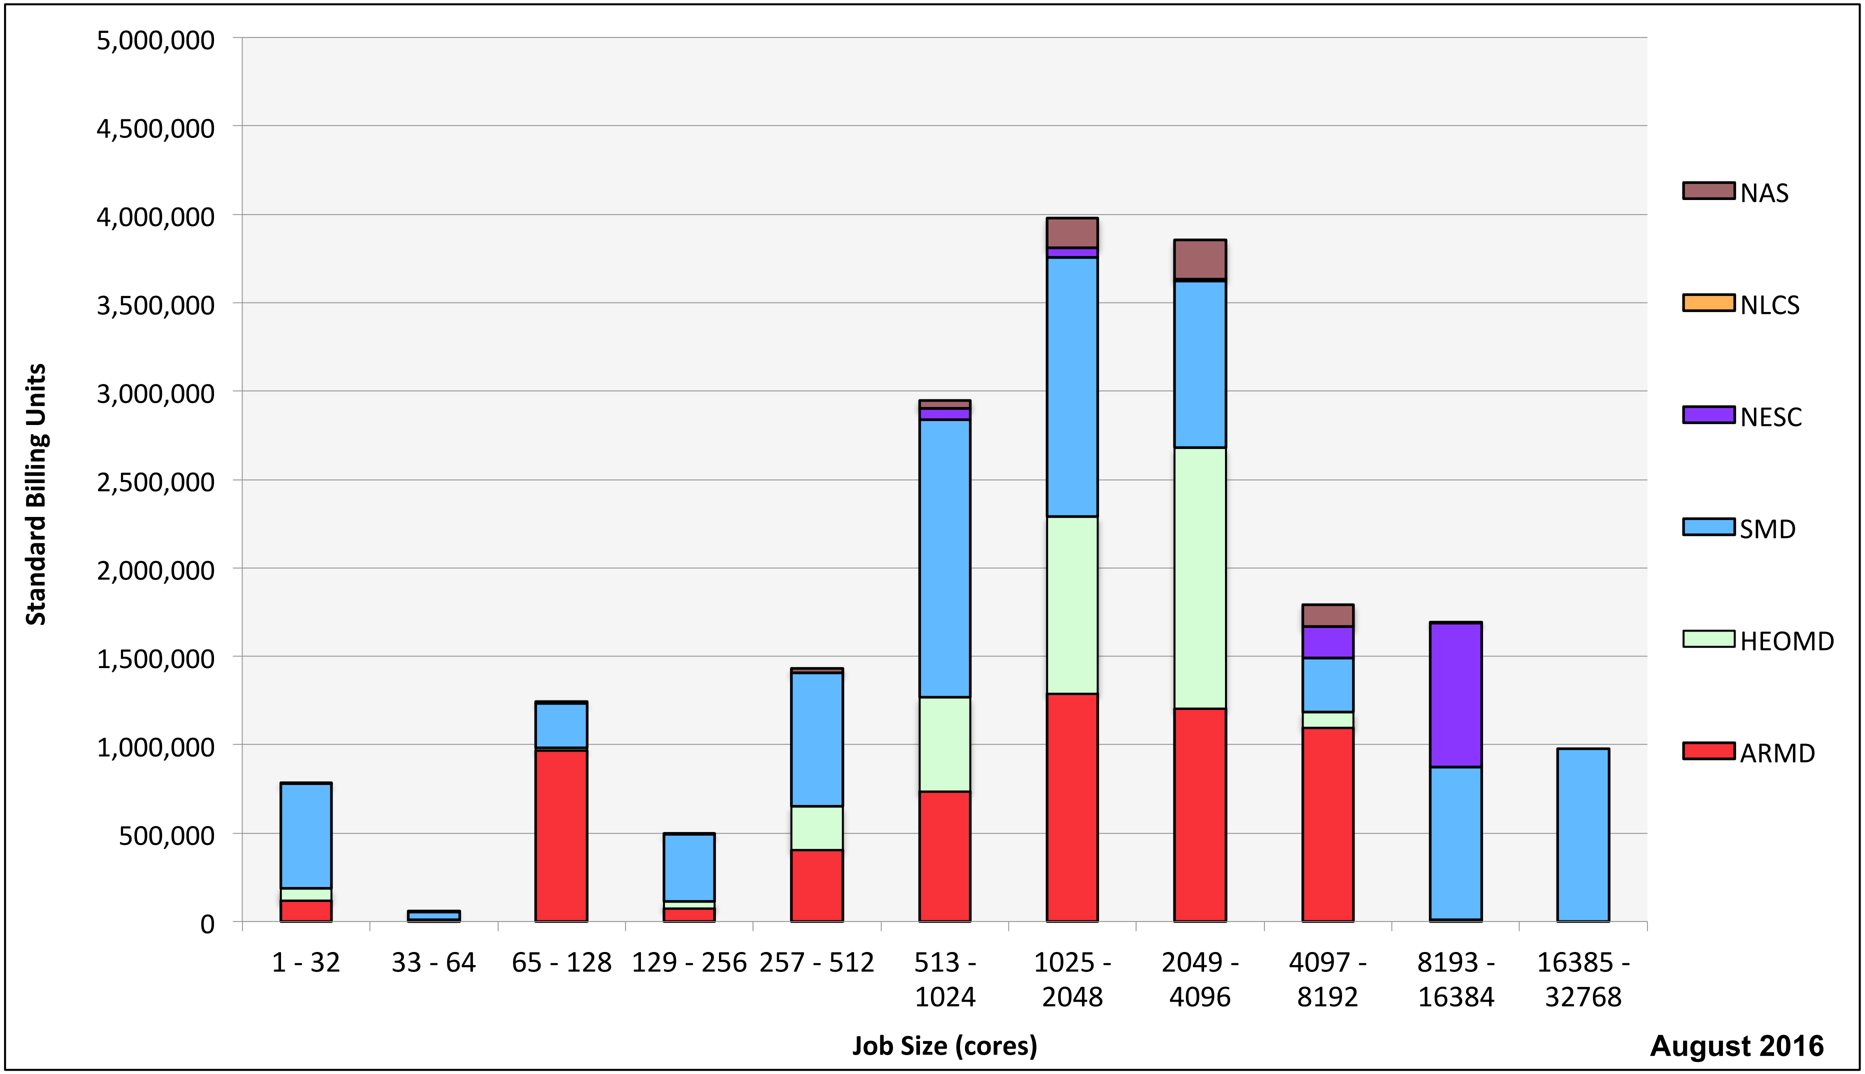

Pleiades: Monthly Utilization by Size and Length

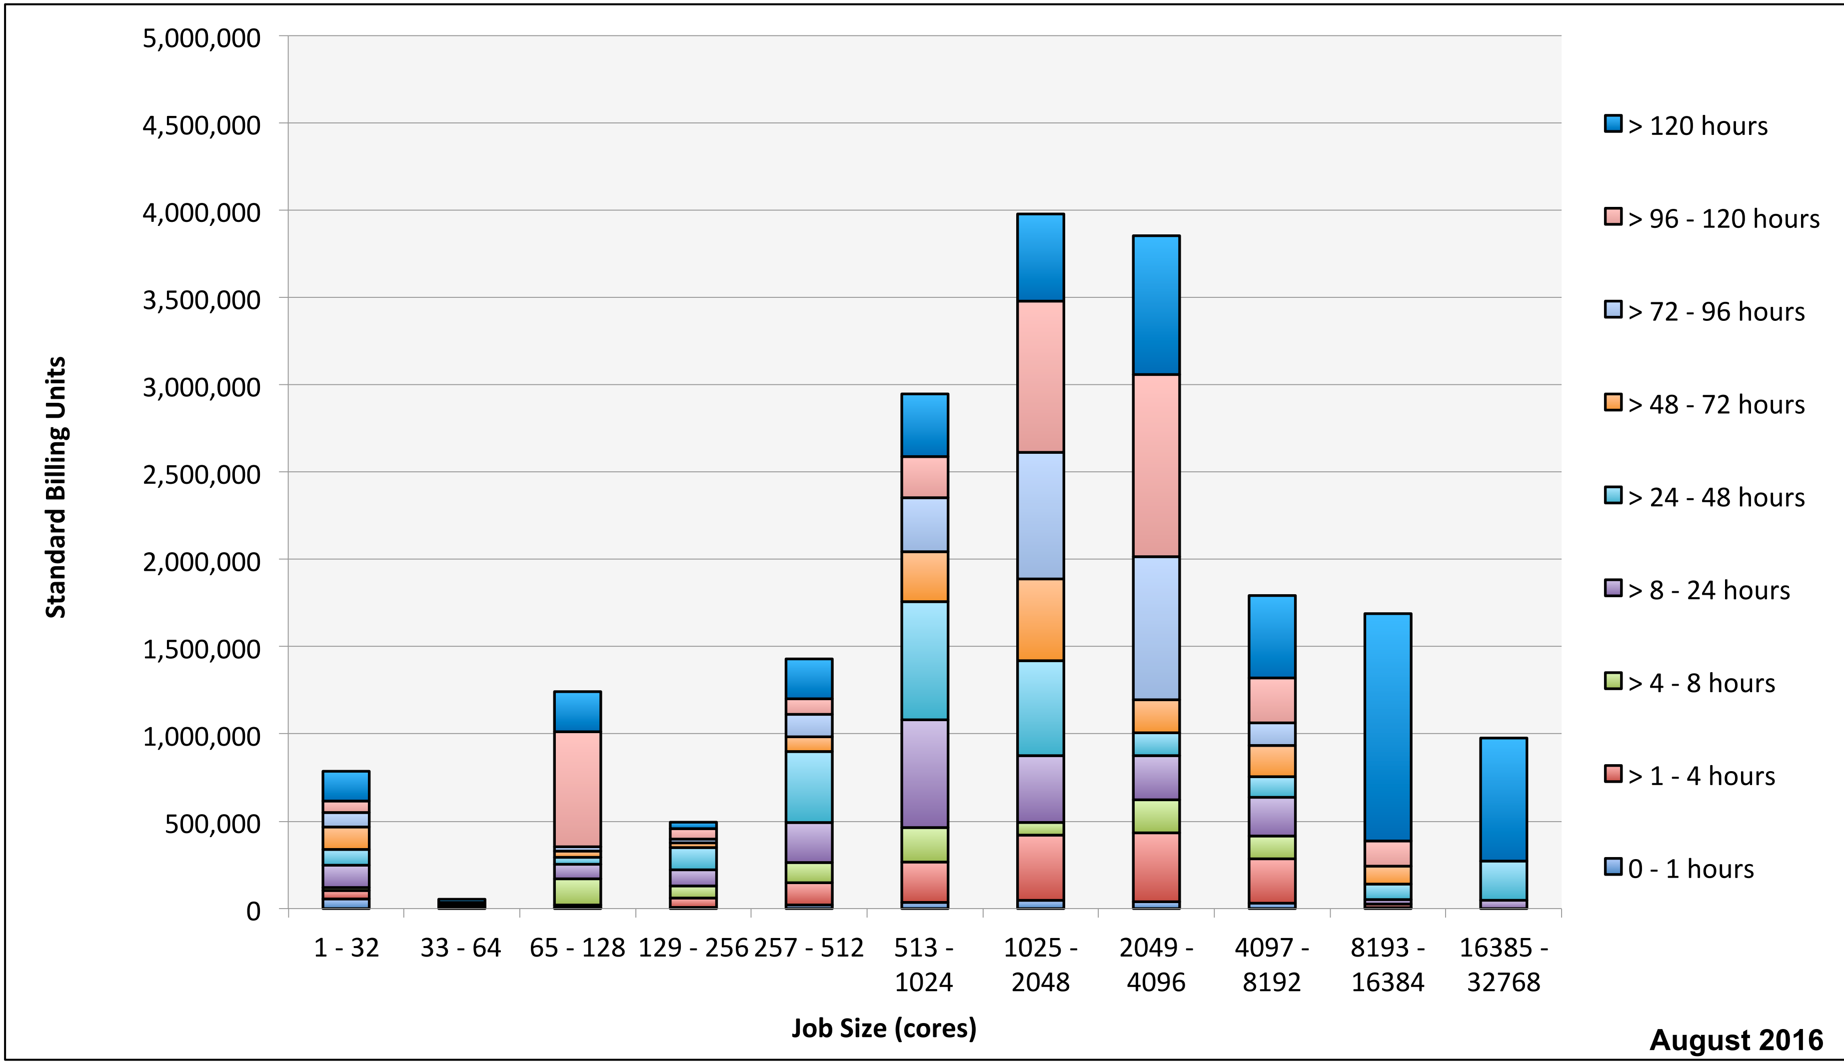

Pleiades: Average Time to Clear All Jobs

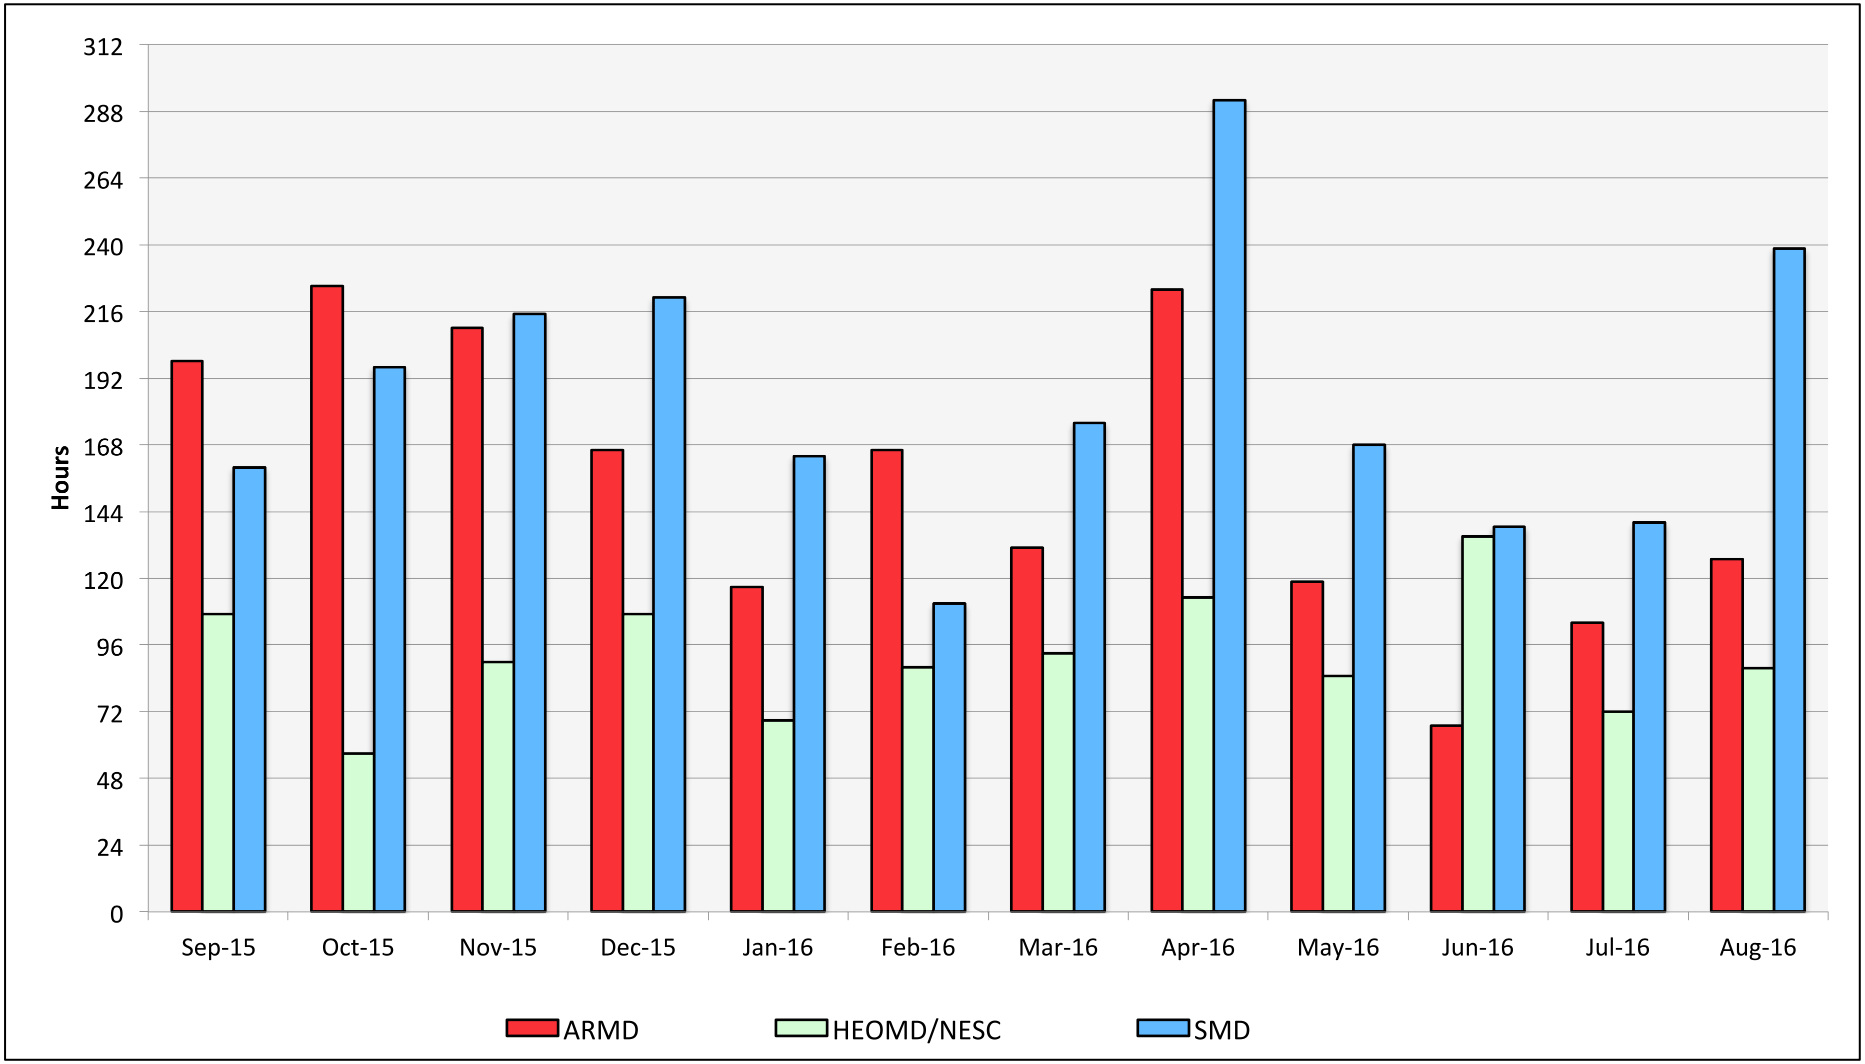

Pleiades:
Average Expansion Factor

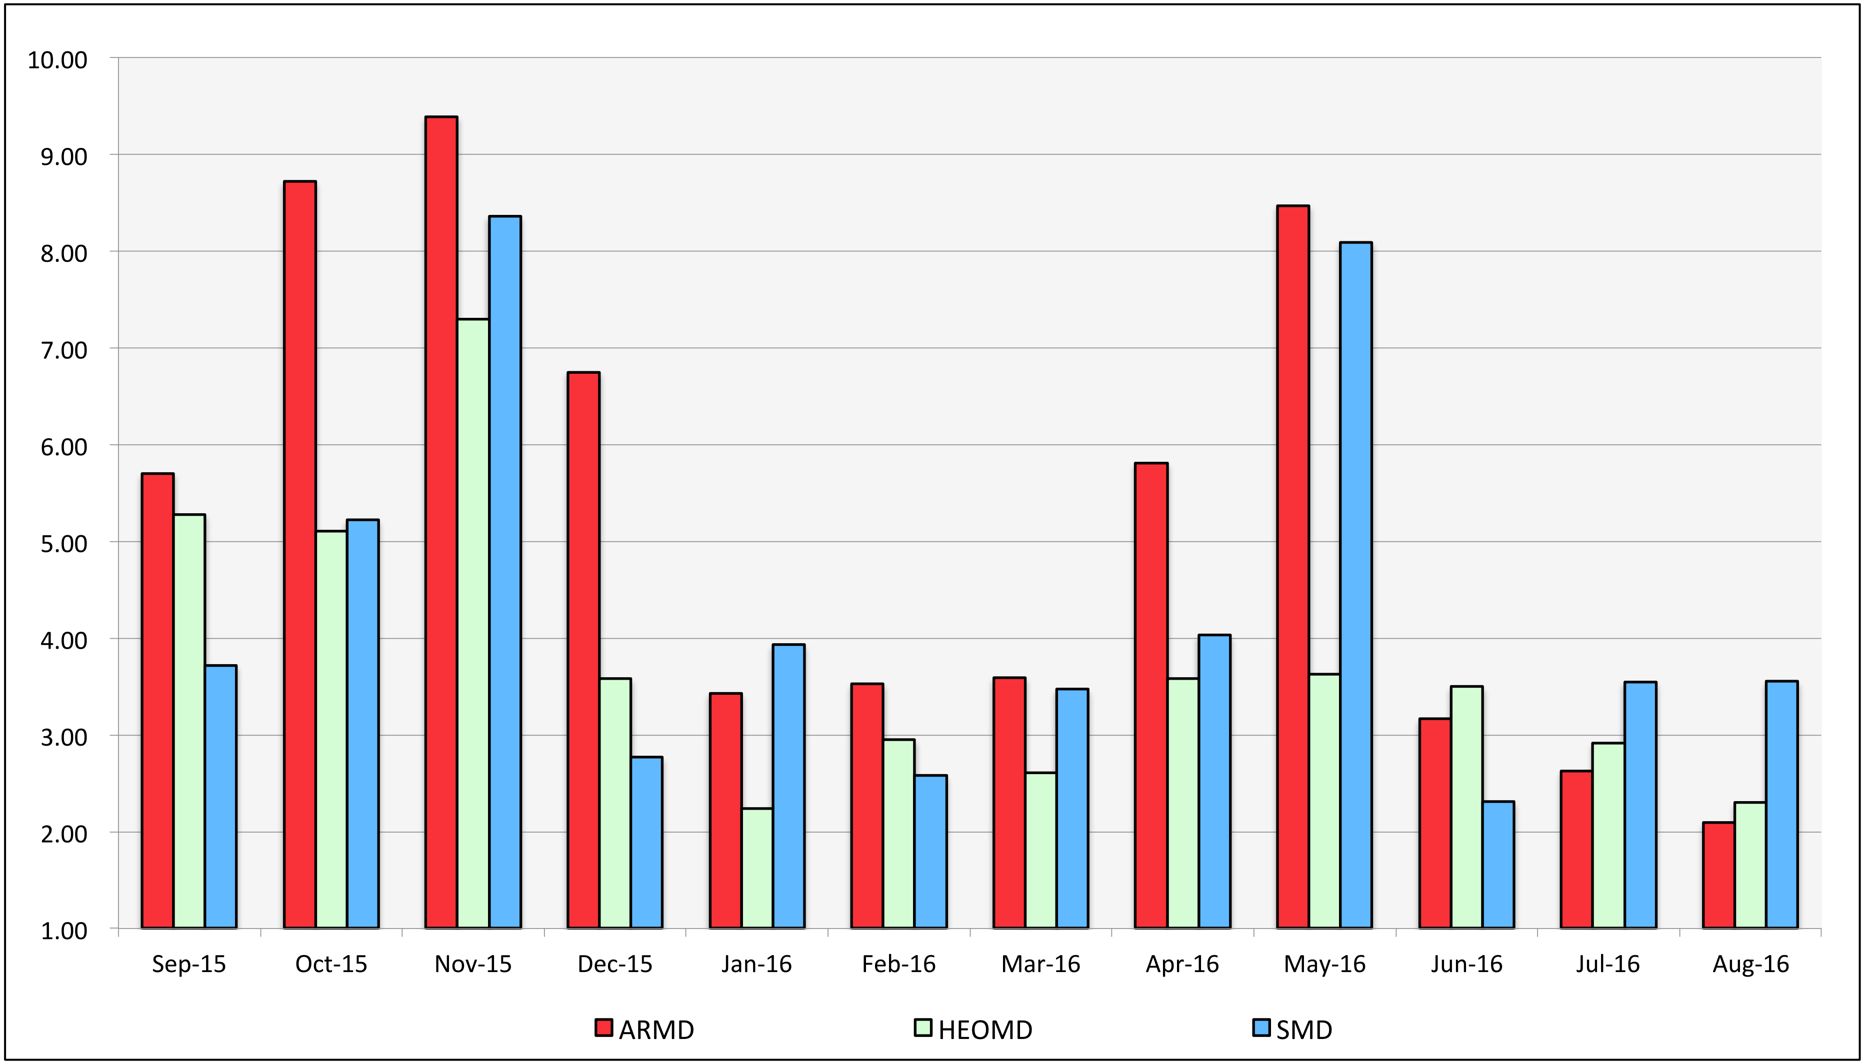

Endeavour: SBUs Reported, Normalized to 30-Day Month

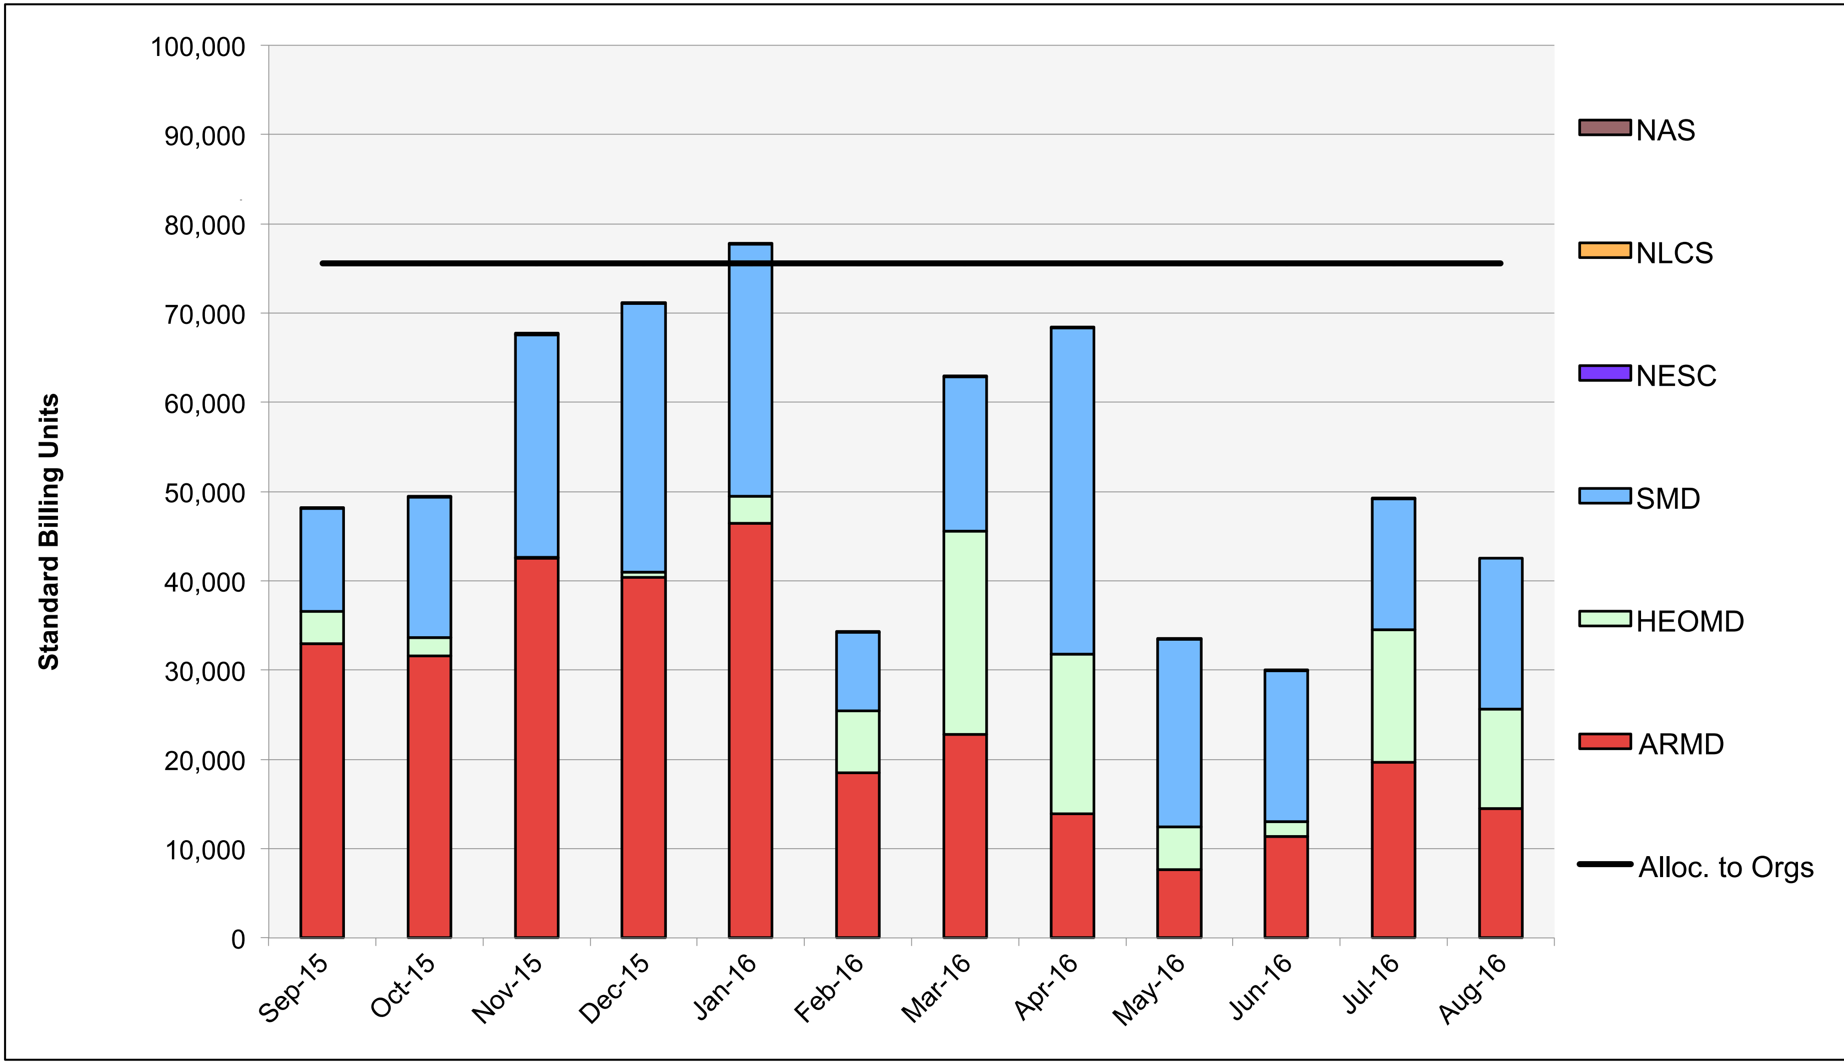

Endeavour: Monthly Utilization by Job Length

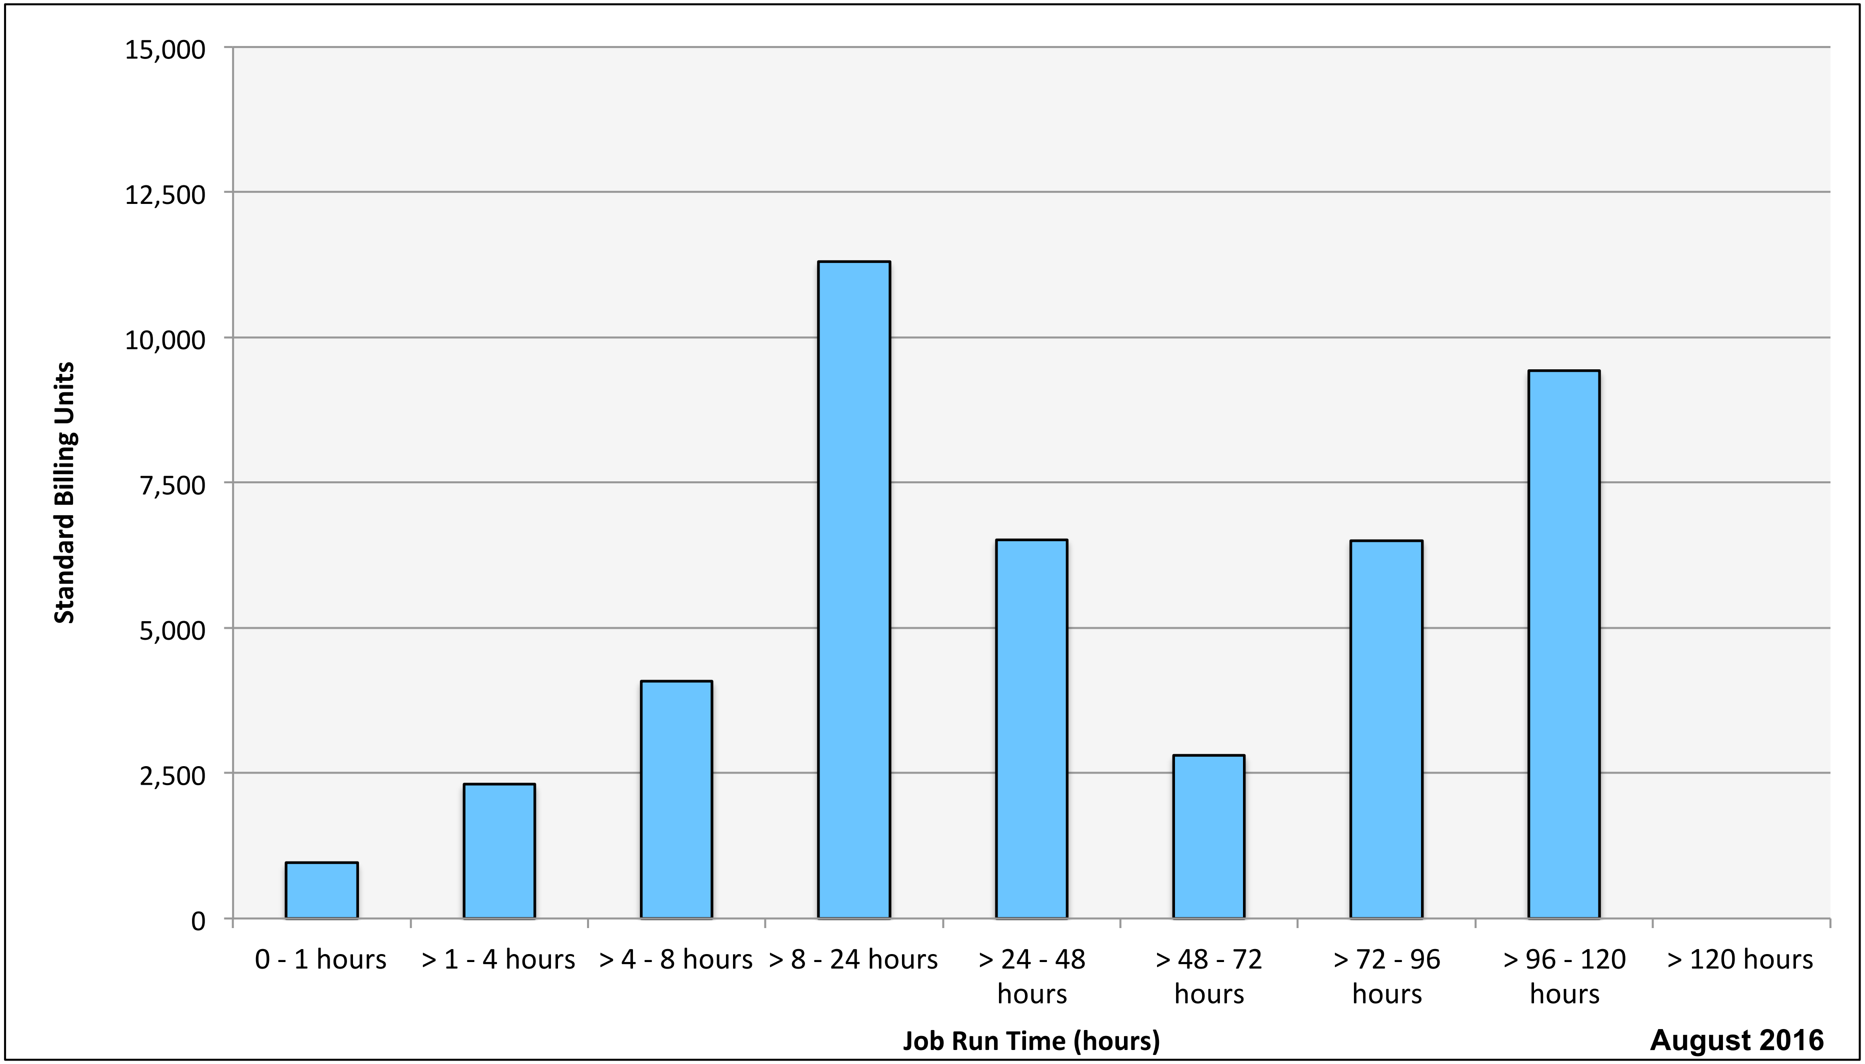

Endeavour: Monthly Utilization by Size and Mission

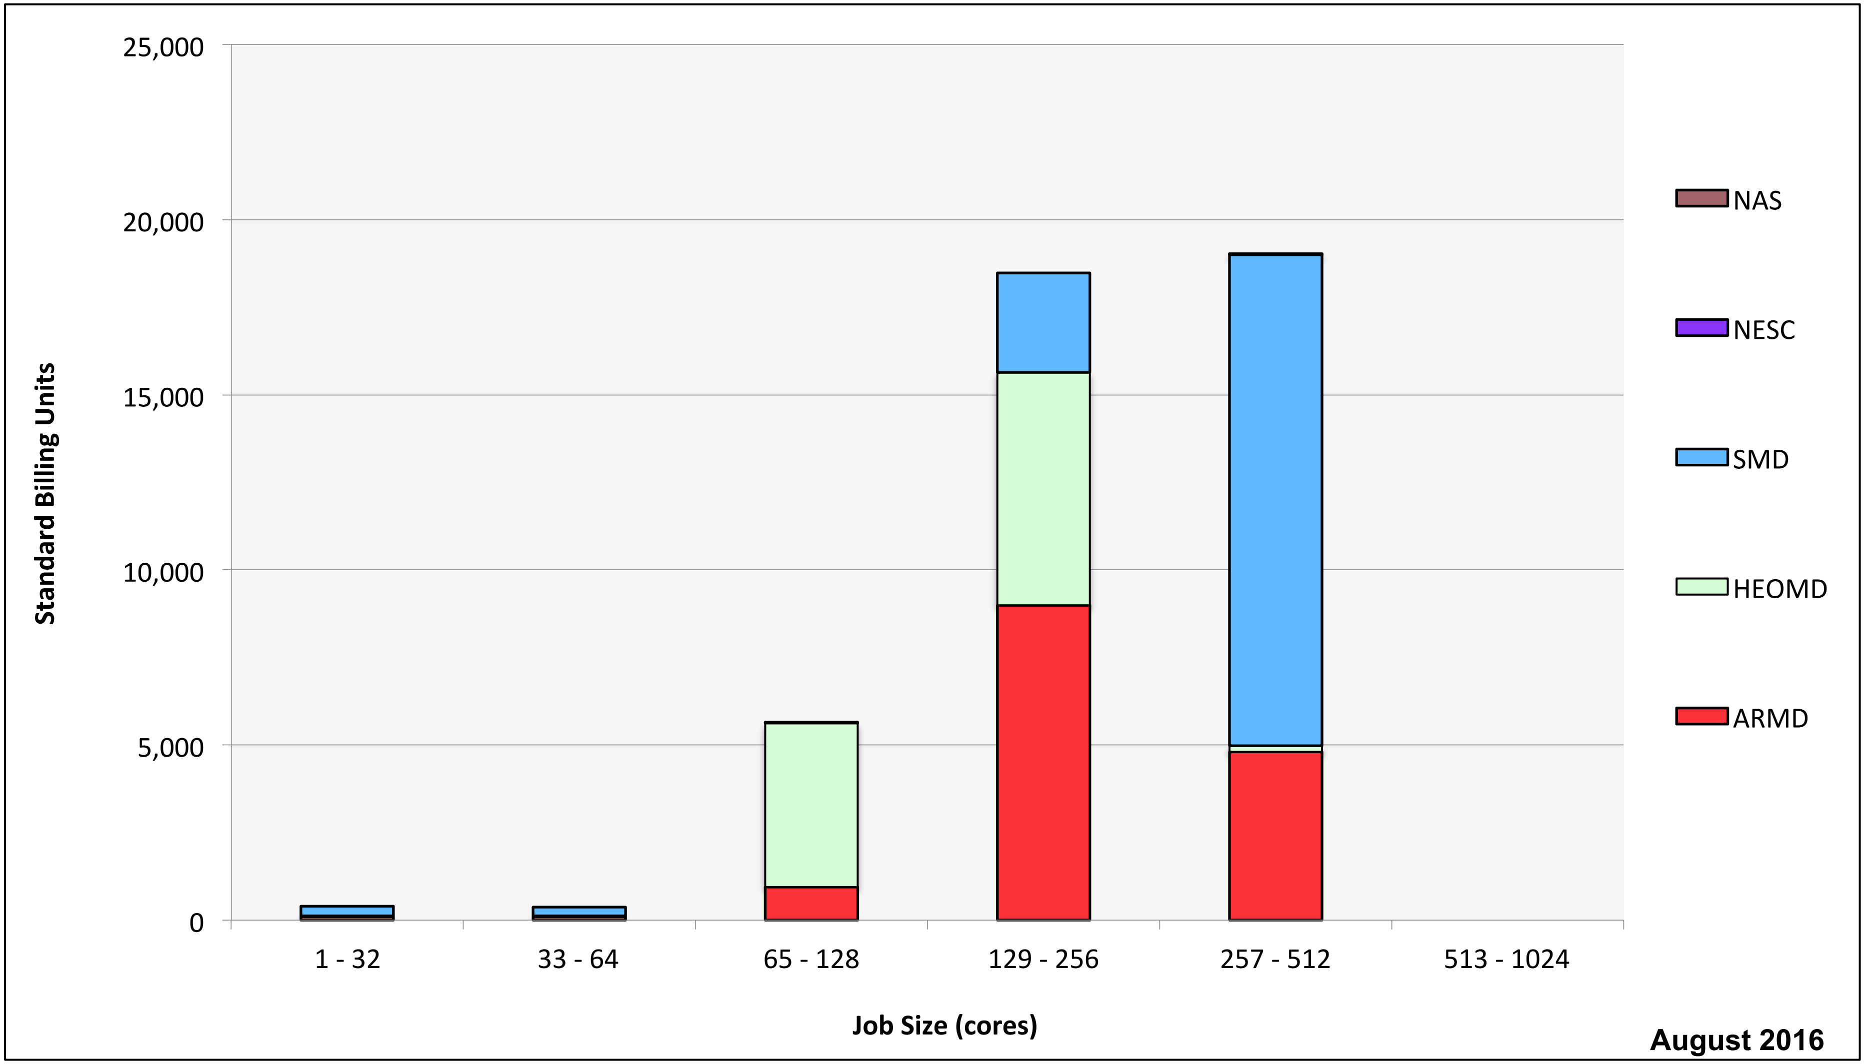

Endeavour: Monthly Utilization by Size and Length

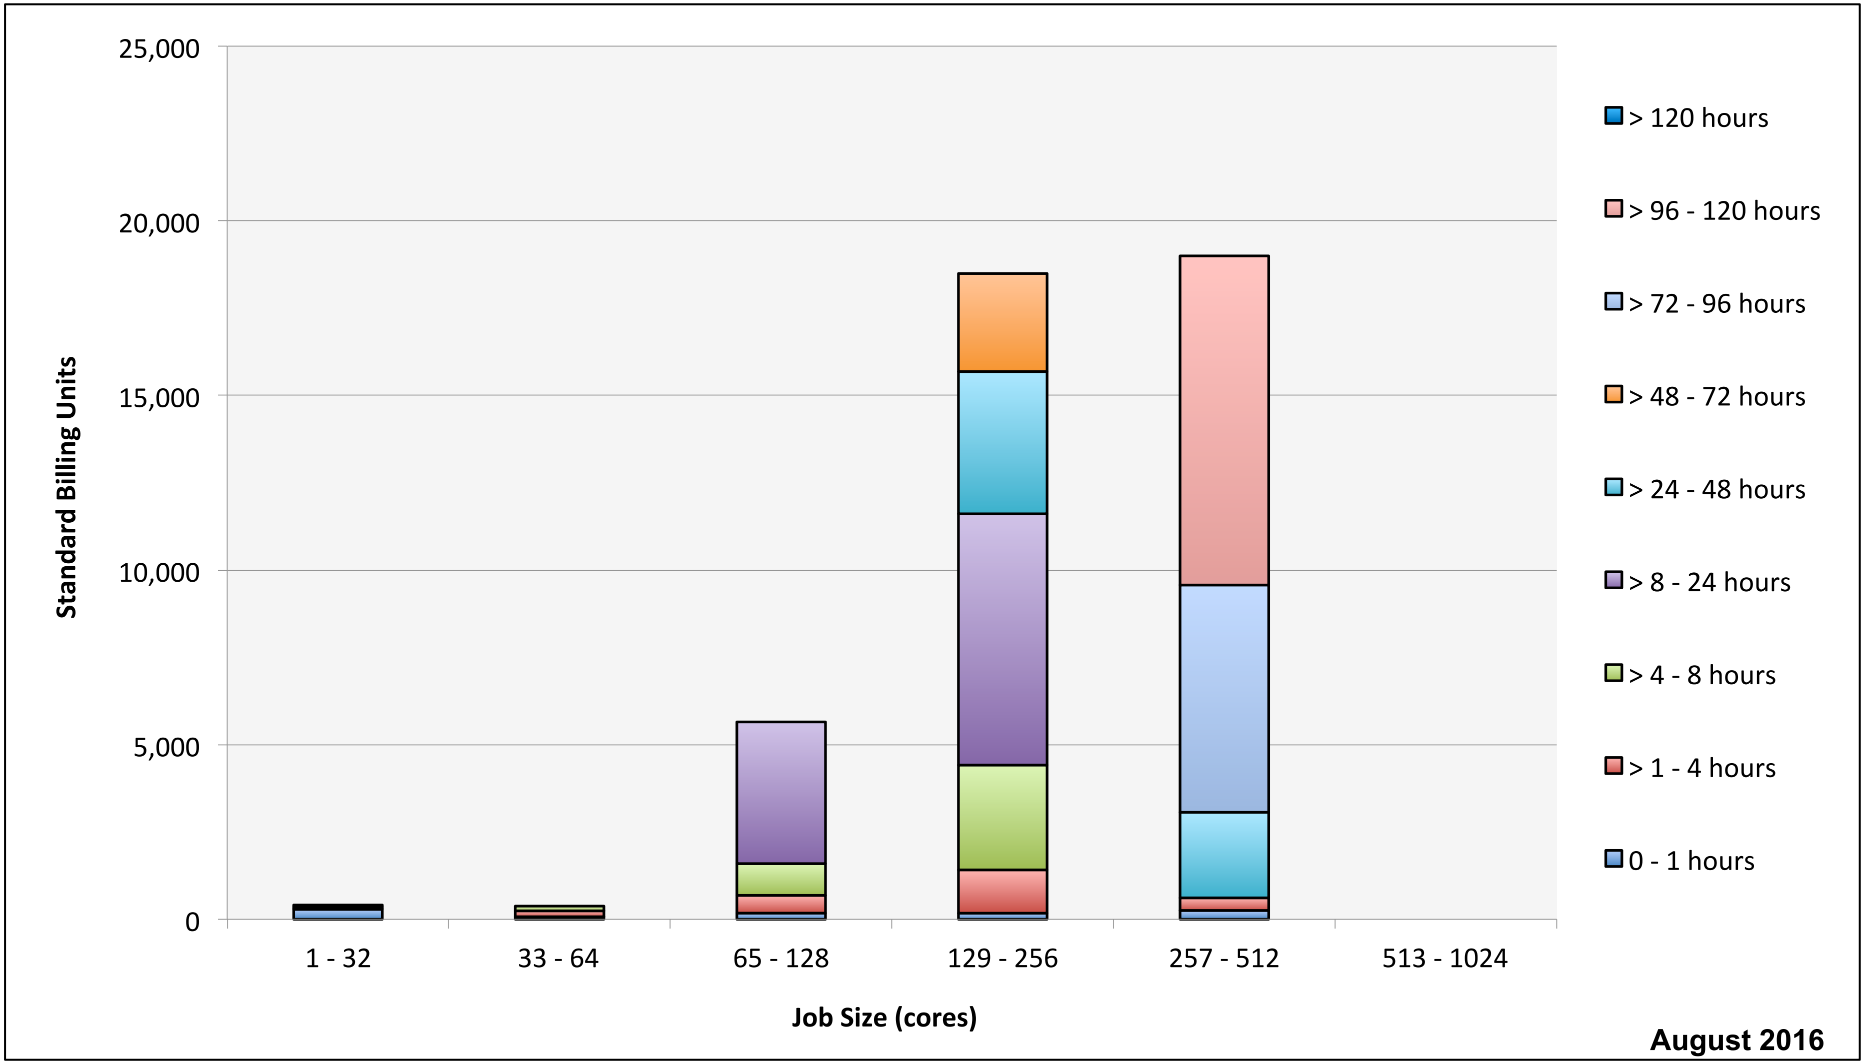

Endeavour:
Average Time to Clear All Jobs

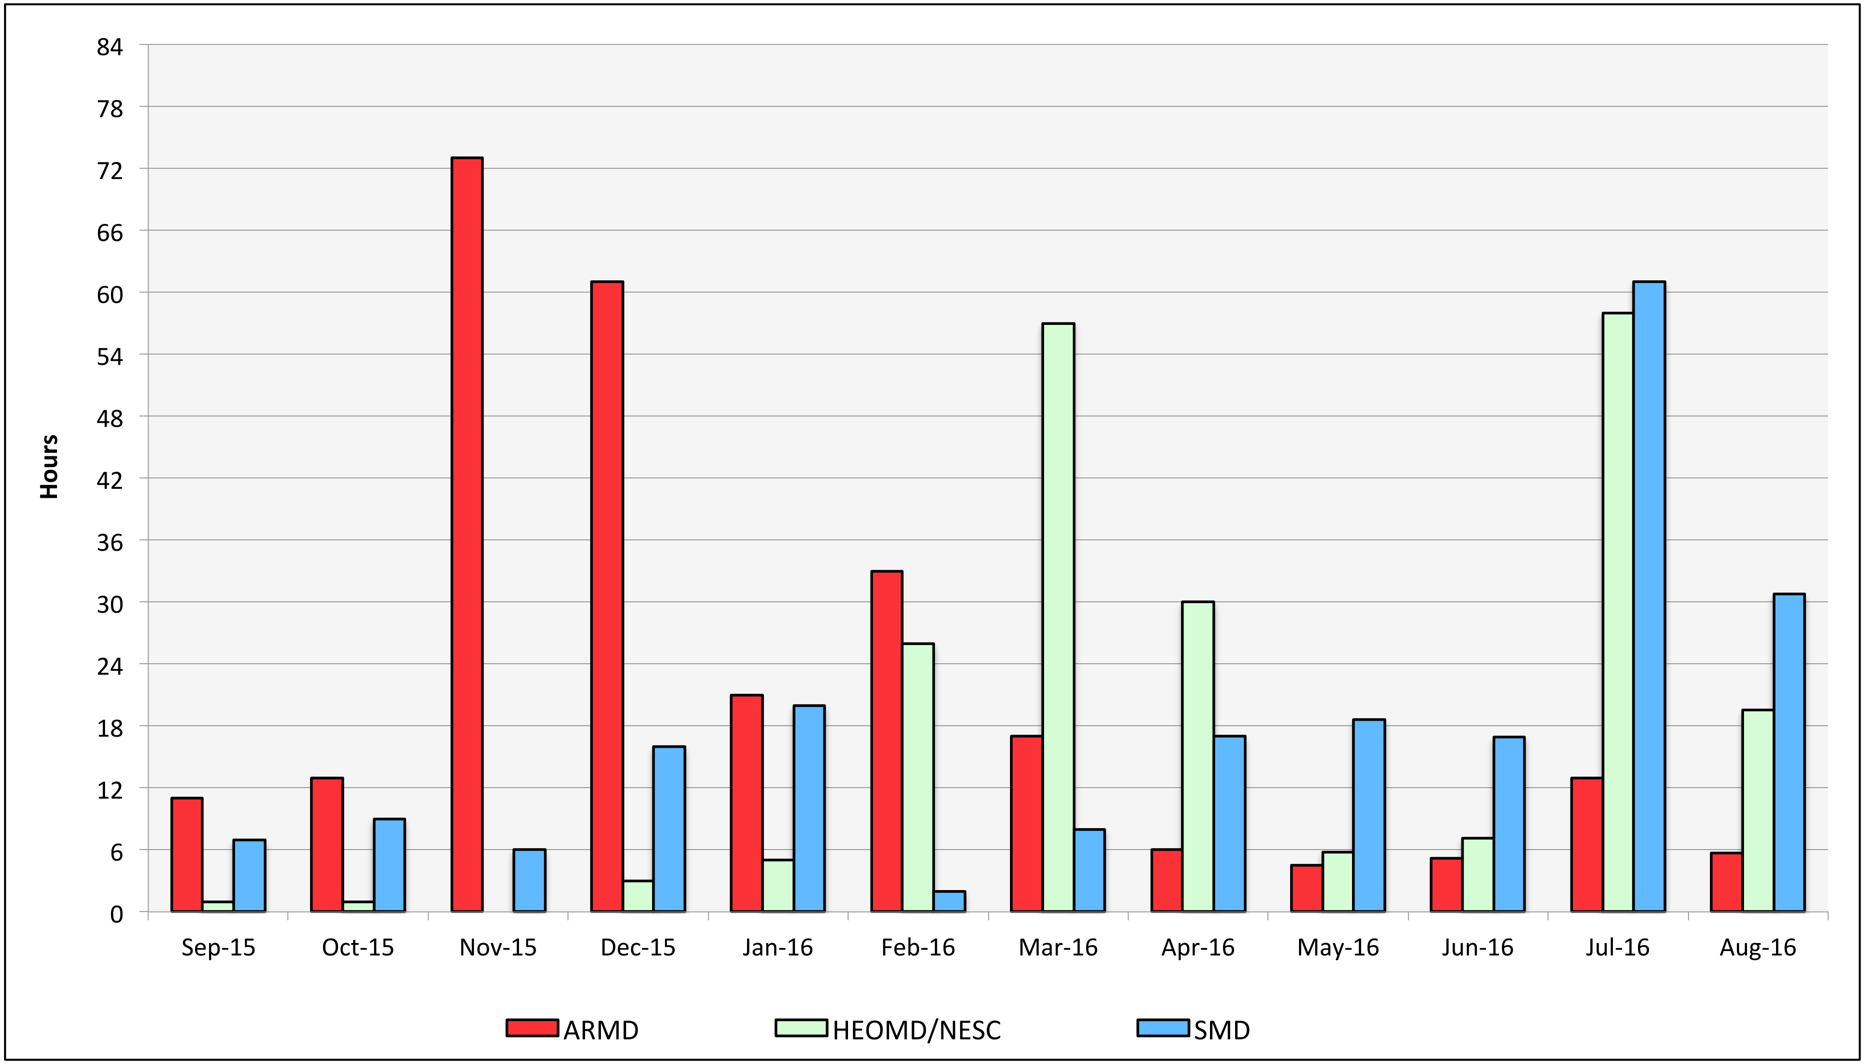

Endeavour:
Average Expansion Factor

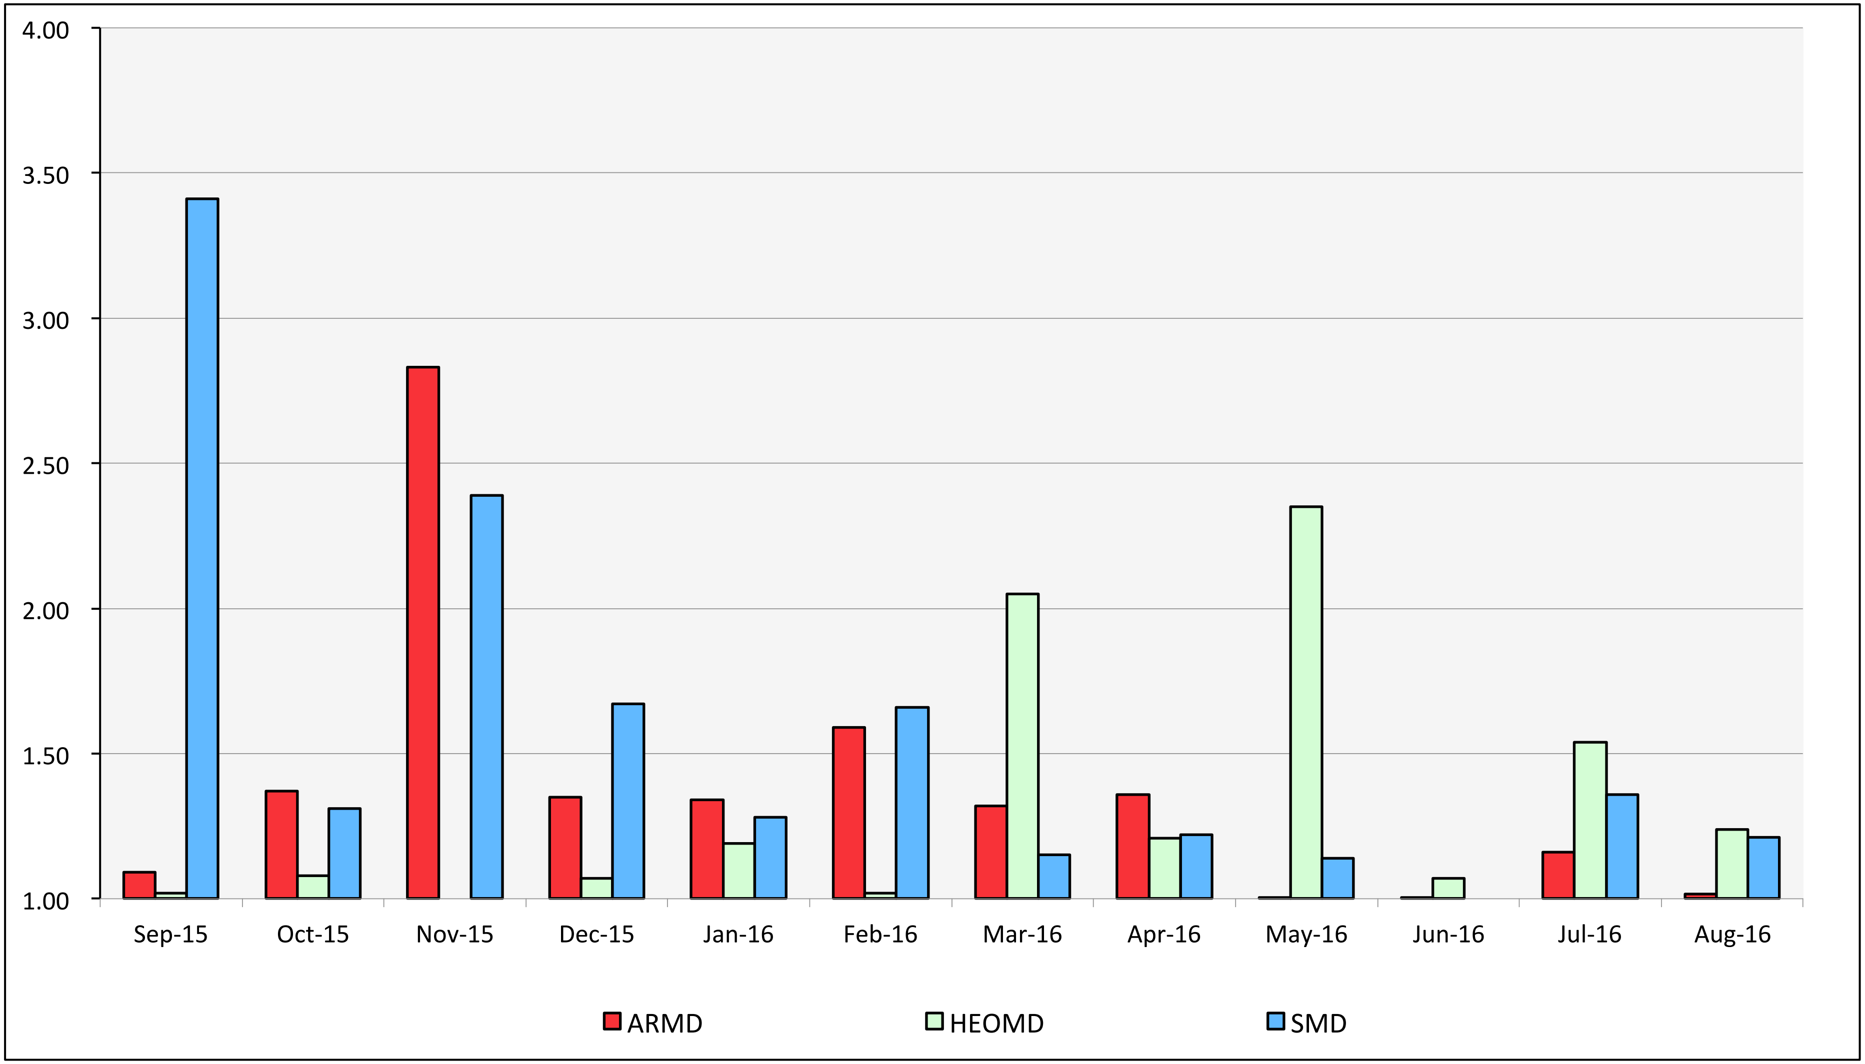

Merope: SBUs Reported, Normalized to 30-Day Month

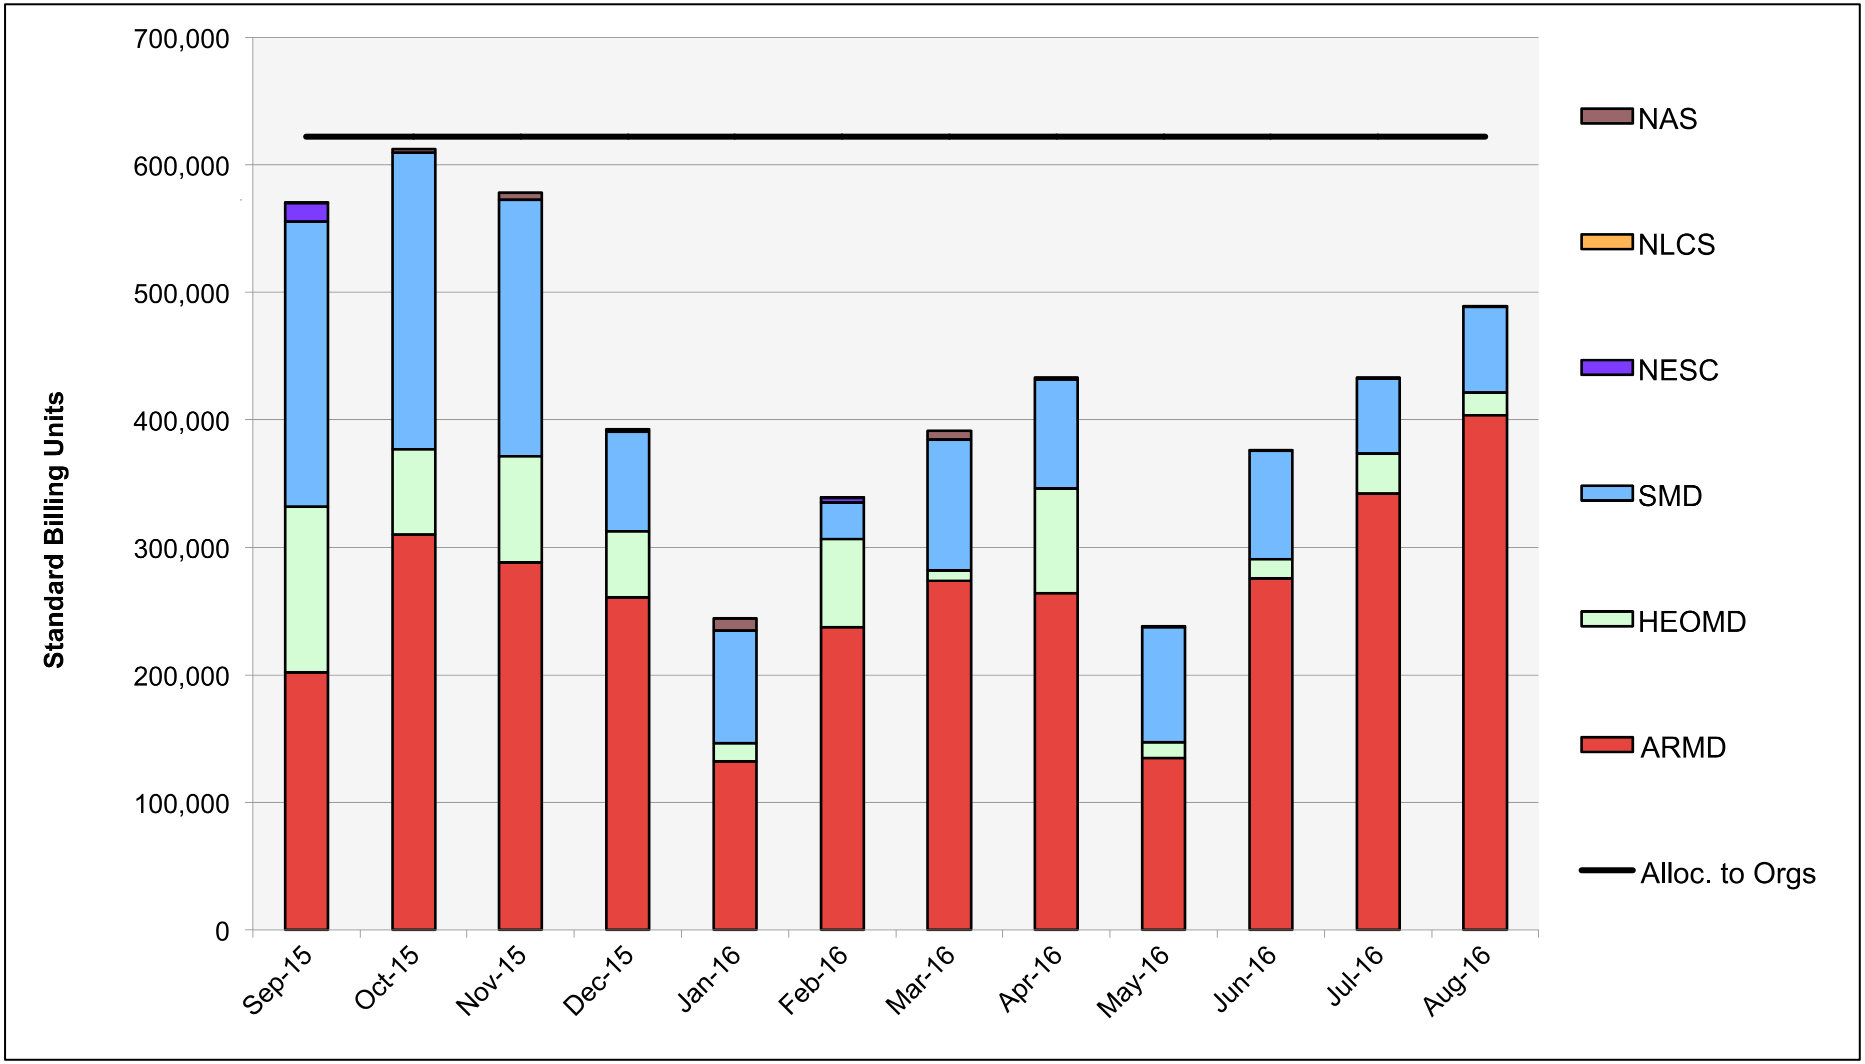

Merope:
Monthly Utilization by Job Length

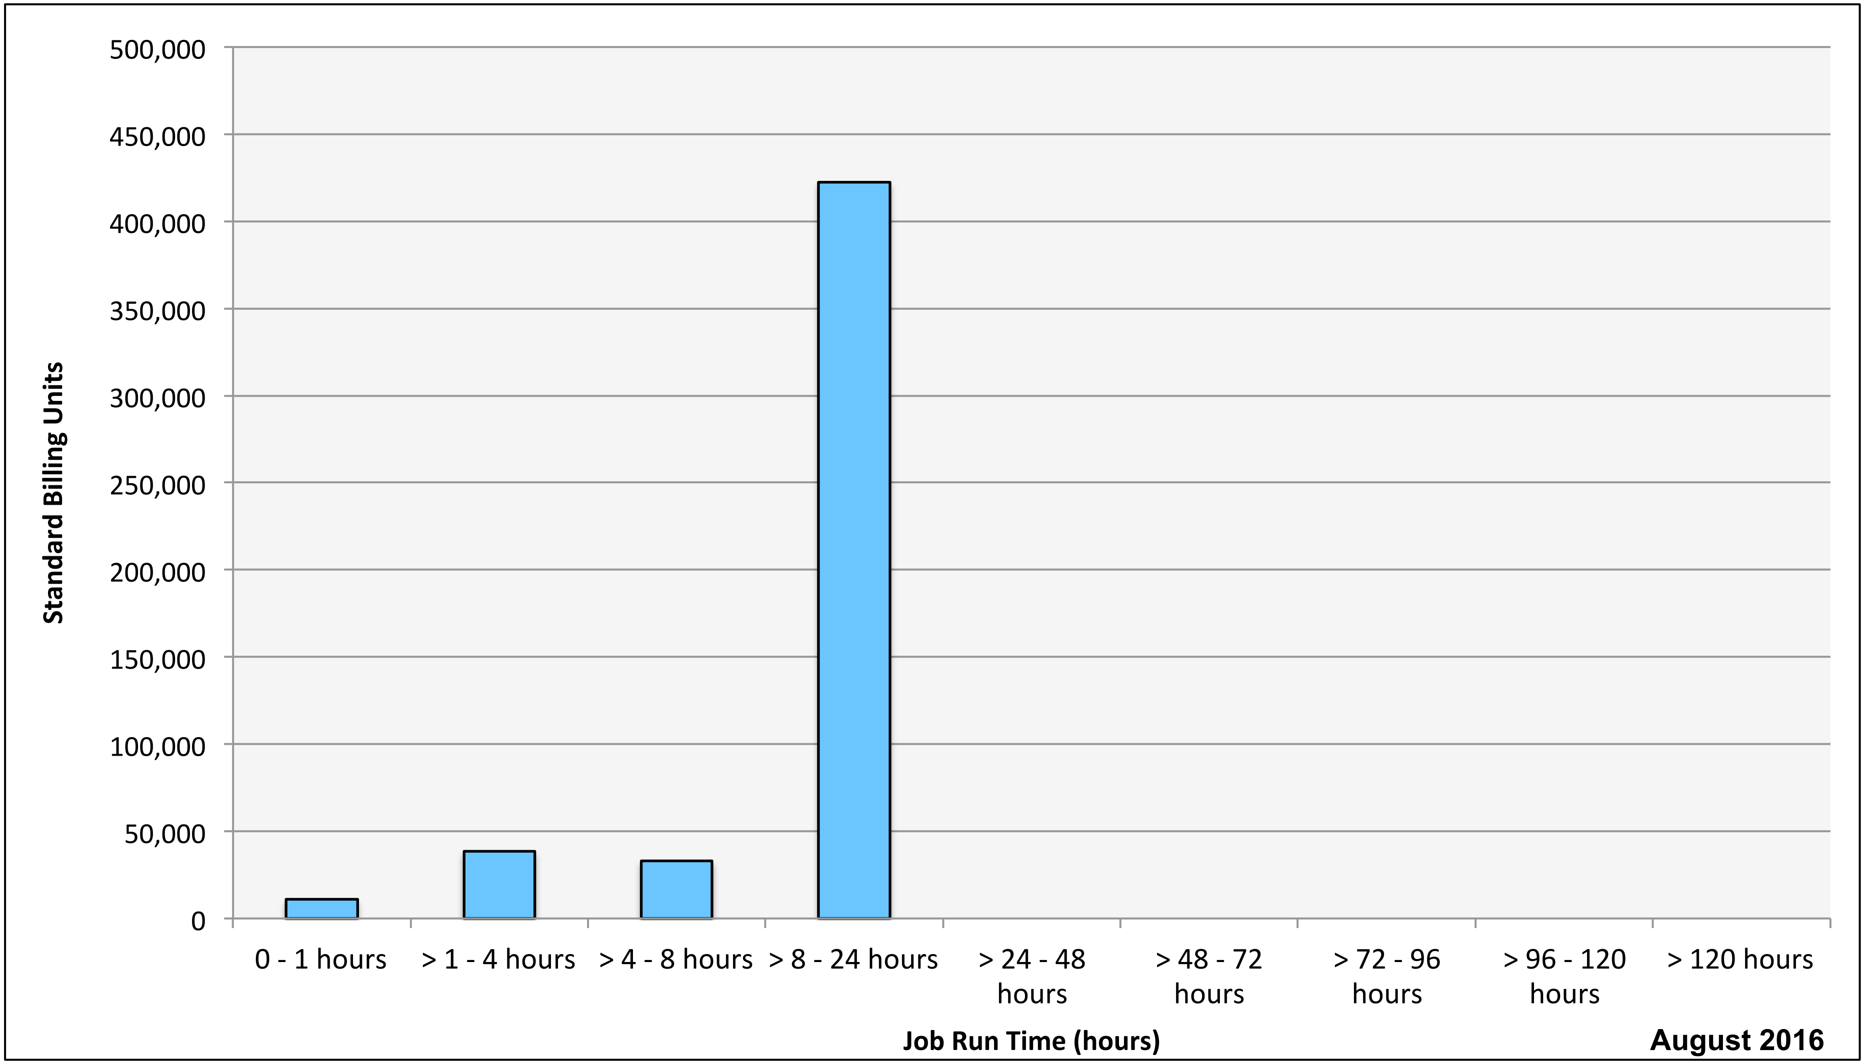

Merope: Monthly Utilization by Size and Mission

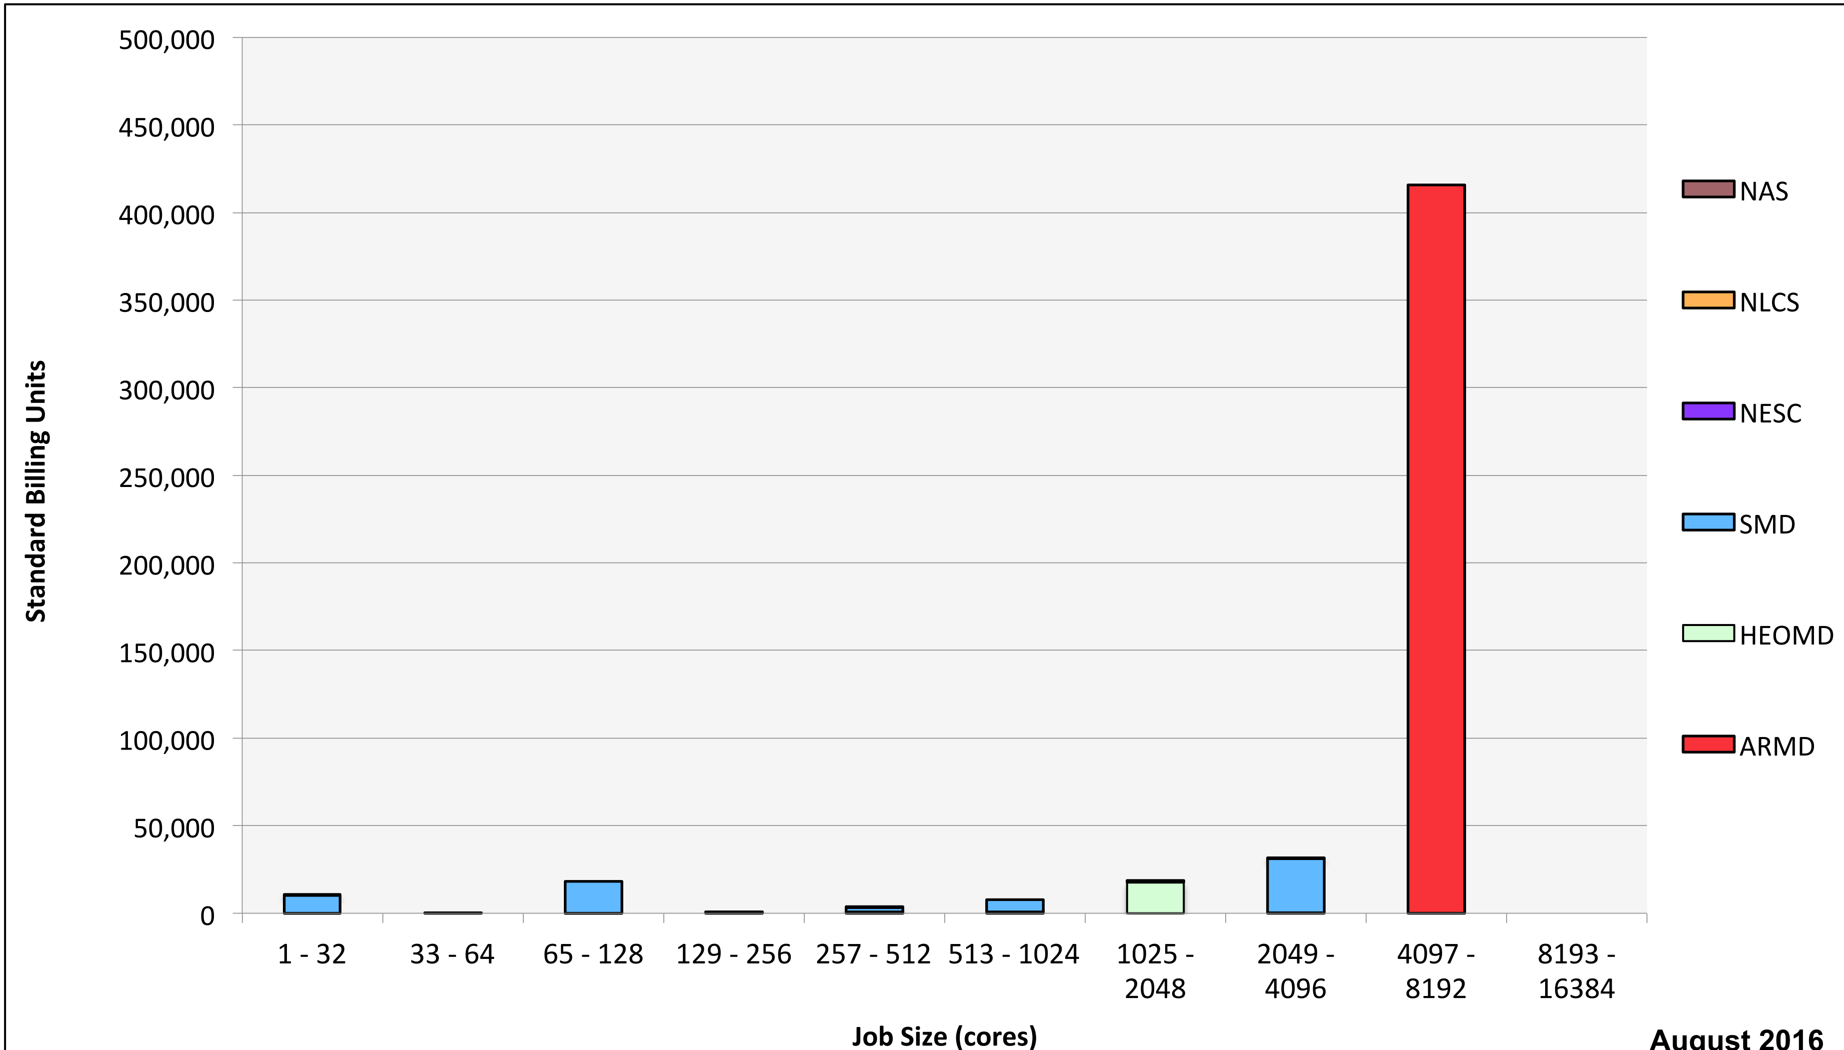

Merope:
Monthly Utilization by Size and Length

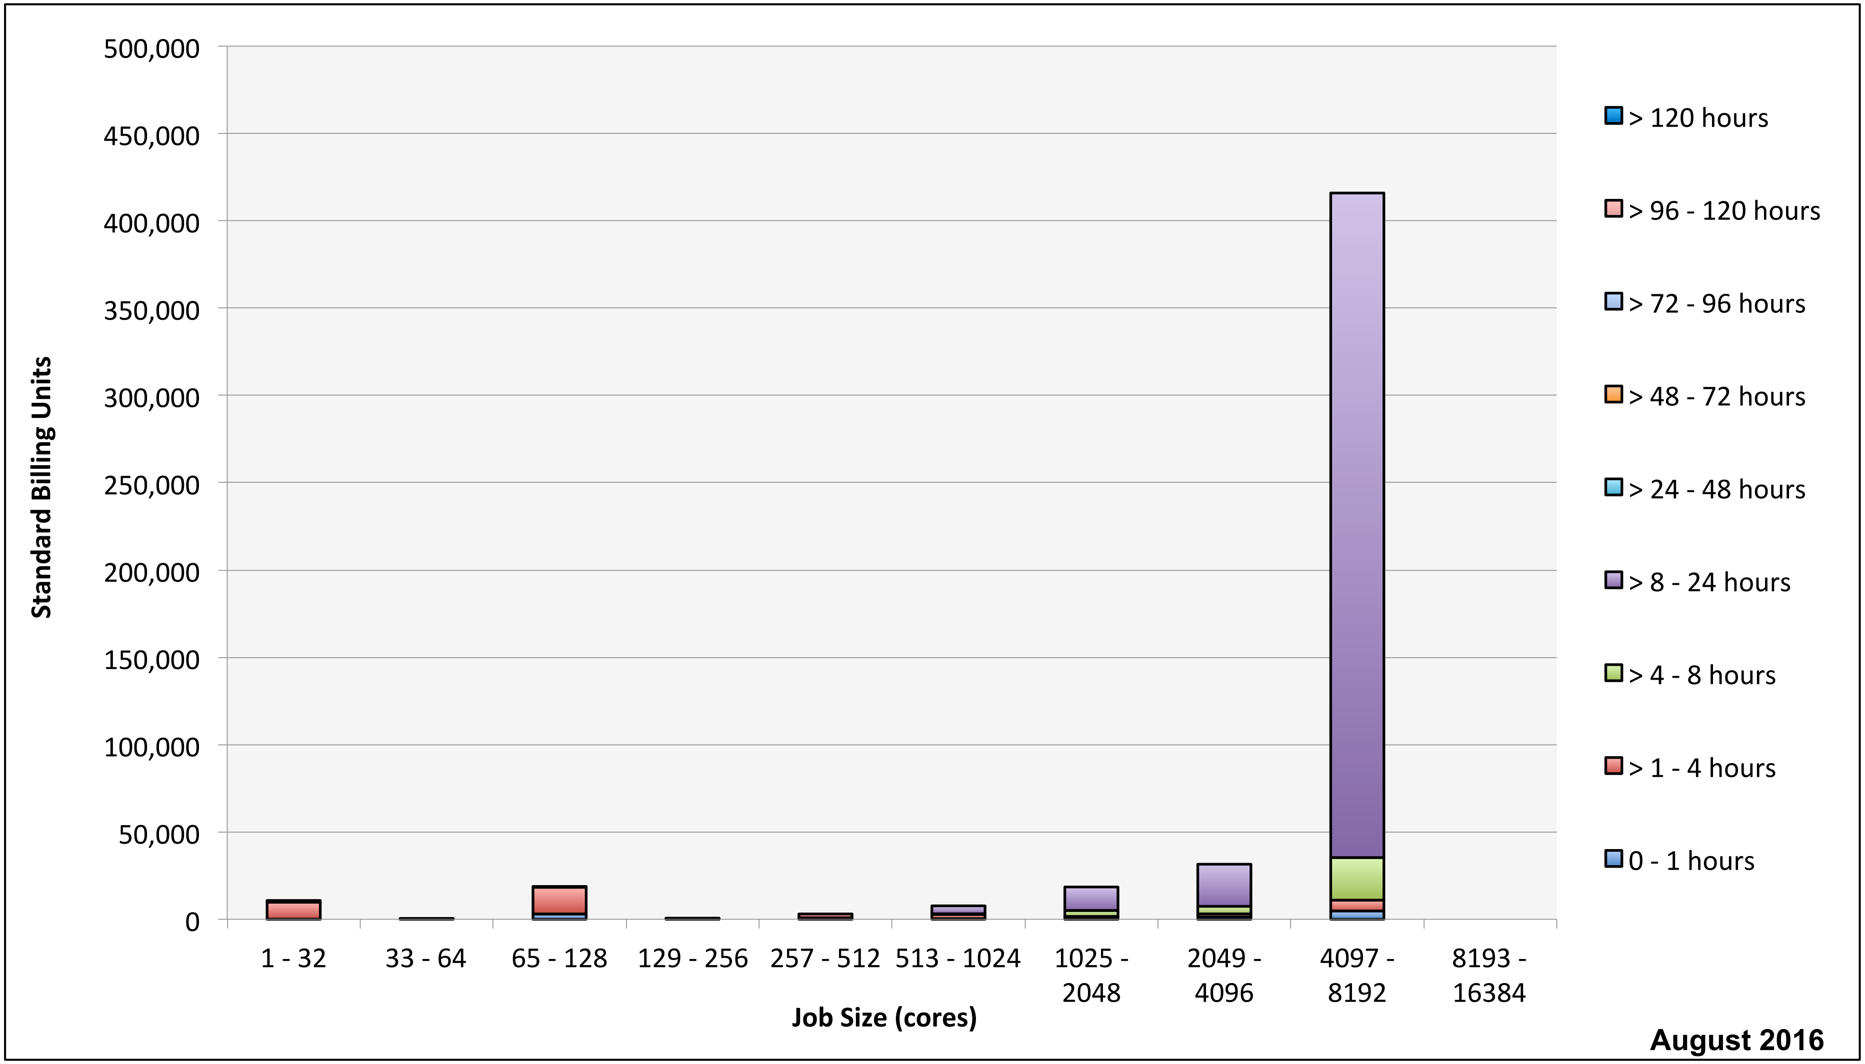

Merope:
Average Time to Clear All Jobs

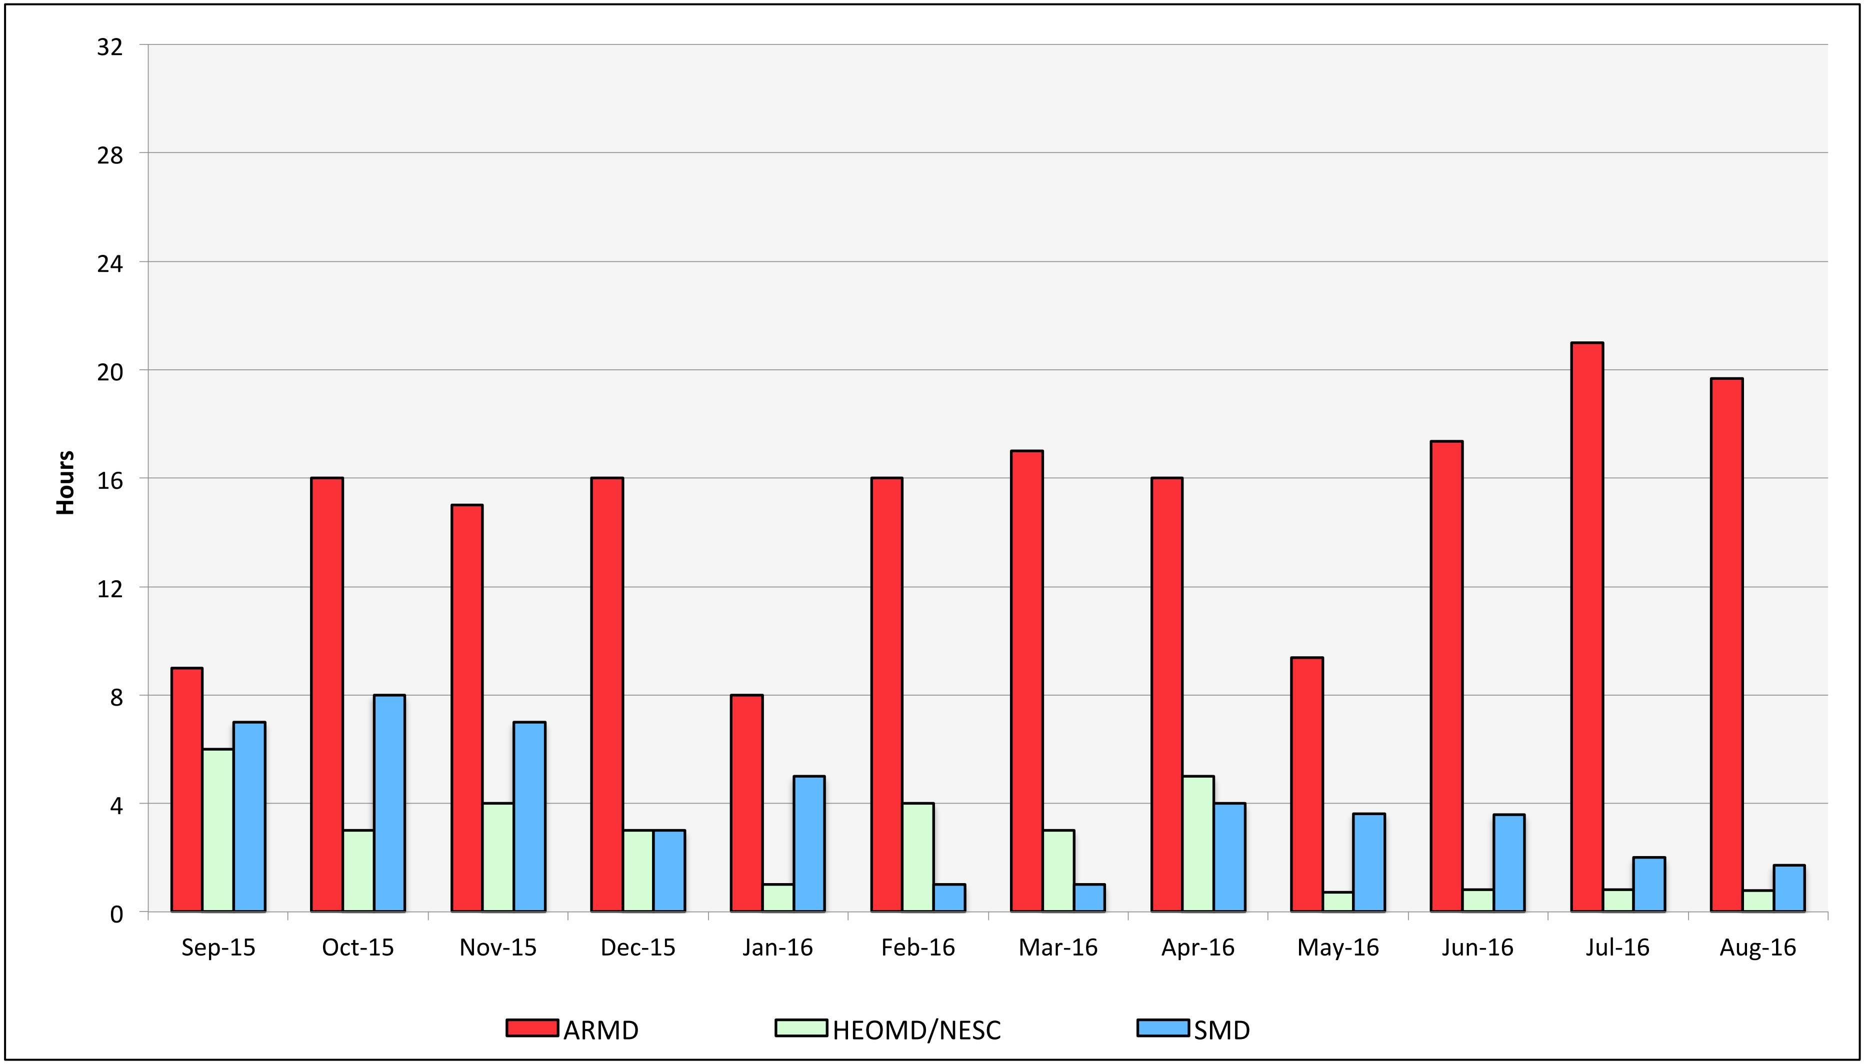

Merope:
Average Expansion Factor

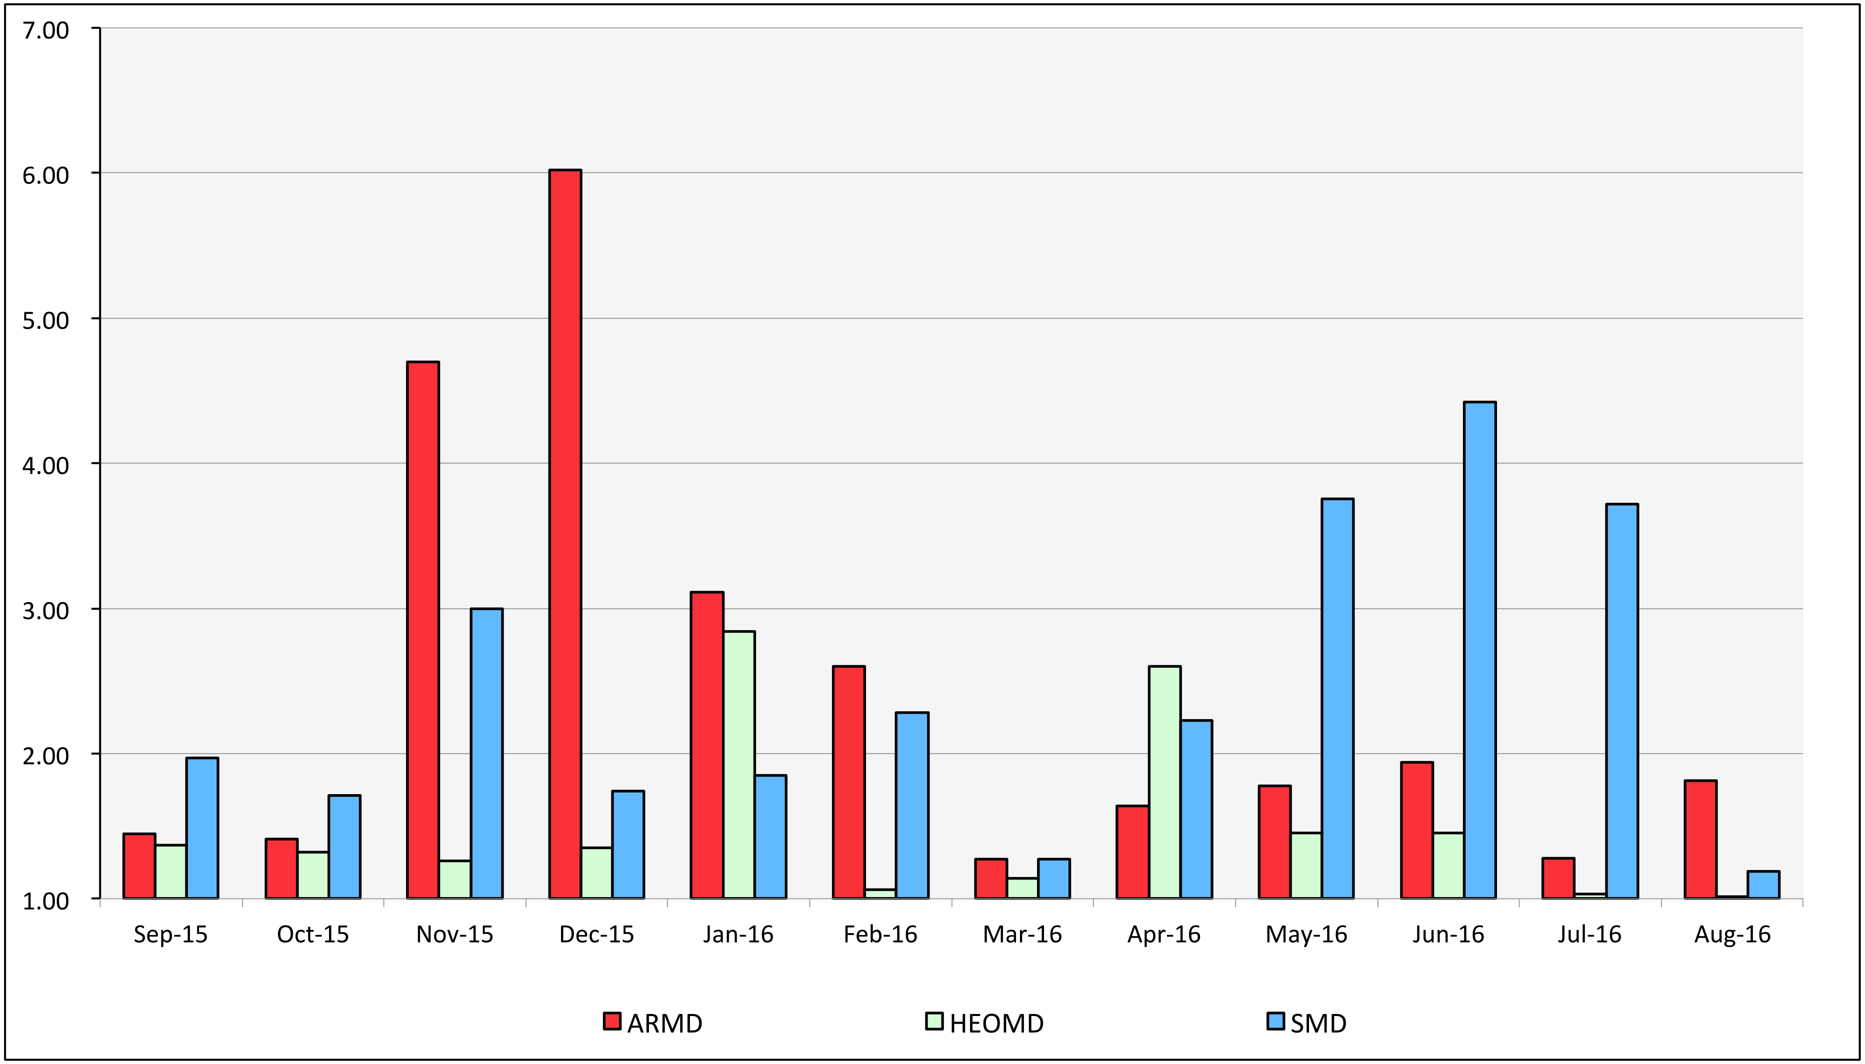

I/O Optimization: ATHENA C++ Application